WL-TR-96-3039

DRA/CHS/HS3/TR95001/01

THE HUMAN-ELECTRONIC CREW:
CAN WE TRUST THE TEAM?

Proceedings of the 3rd International Workshop on
Human-Computer Teamwork

Cambridge, United Kingdom, 27-30 September 1994

Edited by:

Robert M Taylor
DRA Centre for Human
Sciences

John Reising
USAF Wright
Laboratory

DECEMBER 1995

FINAL REPORT FOR 9/27/94--9/30/94

Approved for public release; distribution unlimited

19960524 018

Defence Research Agency
Centre for Human Sciences
Farnborough, Hampshire GU14 6SZ

Published January 1995

FLIGHT DYNAMICS DIRECTORATE
WRIGHT LABORATORY
AIR FORCE MATERIEL COMMAND
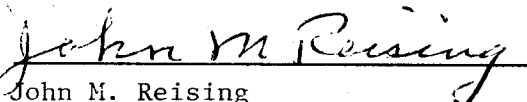WRIGHT-PATTERSON AIR FORCE BASE, OH 45433 –7562
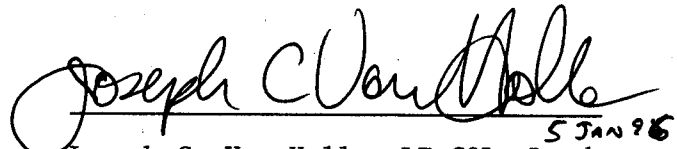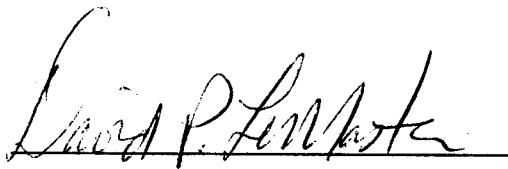
DTIC QUALITY INSPECTED 1

## NOTICE

When Government drawings, specifications, or other data are used for any purpose other than in connection with a definitely Government-related procurement, the United States Government incurs no responsibility or any obligation whatsoever. The fact that the government may have formulated or in any way supplied the said drawings, specifications, or other data, is not to be regarded by implication, or otherwise in any manner construed, as licensing the holder, or any other person or corporation; or as conveying any rights or permission to manufacture, use, or sell any patented invention that may in any way be related thereto.

This report is releasable to the National Technical Information Service (NTIS). At NTIS, it will be available to the general public, including foreign nations.

The technical report has been reviewed and is approved for publication.

John M. Reising
Engineering Psychologist

Joseph C. Von Holle, LT COL., Leader
Advanced Cockpits TTIPT

5 Jan 96

David P. LeMaster, Chief
Flight Control Division
Flight Dynamics Directorate

# REPORT DOCUMENTATION PAGE

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

| 1. AGENCY USE ONLY (Leave blank) | 2. REPORT DATE 19 Dec 95 | 3. REPORT TYPE AND DATES COVERED Conference Proc, 27 to 30 Sep 94 |
|---|---|---|

**4. TITLE AND SUBTITLE**
The Human-Electronic Crew: Can We Trust The Team?

**5. FUNDING NUMBERS**
PE 62201F
PR 2403
TA 04
WU 86

**6. AUTHOR(S)**
Robert M. Taylor and John Reising

**7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)**
FLIGHT DYNAMICS DIRECTORATE
WRIGHT LABORATORY
AIR FORCE MATERIEL COMMAND
WRIGHT PATTERSON AIR FORCE BASE OH 45324-7562

**8. PERFORMING ORGANIZATION REPORT NUMBER**
WL-TR-96-3039

**9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)**
FLIGHT DYNAMICS DIRECTORATE
WRIGHT LABORATORY
AIR FORCE MATERIEL COMMAND
WRIGHT PATTERSON AIR FORCE BASE, OH 45433-7562

**10. SPONSORING/MONITORING AGENCY REPORT NUMBER**
WL-TR-96-3039

**11. SUPPLEMENTARY NOTES**
Prepared in cooperation with the Royal Air Force and German Air Force.

**12a. DISTRIBUTION/AVAILABILITY STATEMENT**
Approved for public release; Distribution is unlimited

**12b. DISTRIBUTION CODE**

**13. ABSTRACT (Maximum 200 words)**

With the emergence of increasing numbers of aircraft systems involving human operators interacting with intelligent automation, concerns have been raised reguarding the trustworthiness of the Human-Electronic crew teams's decisions. Many of the decisions that the team is required to make occur in an imprecise world in which the judgements may be made based on such vague concepts as high, low, near, or far. The most effective decision aids in this type of enviroment may be those which can interpret inexact data and still achieve sound solutions, such as fuzzy logic decision aiding systems. Essentially, the problem comes down to the level of confidence that higher authorities should have in the decisions, and the resulting actions of the team. The Human-Electronic crew needs to be successfully developed and integrated such that it can function effectively as a trustwortht team in this inexact real world.

**14. SUBJECT TERMS**
Electronic Crewmember
Trustworthiness

**15. NUMBER OF PAGES**
198

**16. PRICE CODE**

| 17. SECURITY CLASSIFICATION OF REPORT | 18. SECURITY CLASSIFICATION OF THIS PAGE | 19. SECURITY CLASSIFICATION OF ABSTRACT | 20. LIMITATION OF ABSTRACT |
|---|---|---|---|
| Unclassified | Unclassified | Unclassified | SAR |

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102

# CONTENTS

# OBJECTIVE

With the emergence of increasing numbers of aircraft systems involving human operators interacting with "intelligent" automation, concerns have been raised regarding the trustworthiness of the Human-Electronic Crew Team's decisions. Many of the decisions that the Team is required to make occur in an imprecise world in which the judgements may be made based on such vague concepts as high, low, near or far. The most effective decision aids in this type of environment may be those which can interpret inexact data and still achieve sound solutions, such as fuzzy logic decision aiding systems. Essentially, the problem comes down to the level of confidence that higher authorities should have in the decisions, and in the resulting actions, of the Team. The Human-Electronic Crew needs to be successfully developed and integrated such that it can function effectively as a trustworthy Team in this inexact, real world. The specific purpose of this workshop was to examine these concerns.

This workshop was a follow-up to two previously successful meetings (1988 and 1990) co-sponsored by the RAF and USAF. It provided a timely forum for experts of several countries to measure progress in this critical technical area. It allowed for the exchange of new ideas, concepts and data relative to hardware and software capabilities that can be included in aircraft system design, to aid the human operator perform the mission. Attendance at the workshop was by invitation only. The numbers of persons attending was restricted to 60. All invited attendees were expected to contribute through active participation in the meeting discussions. It brought together experts representing cockpit design disciplines including hardware and software technologists, as well as human factors specialists and pilots to address such questions as:

> (1) Do current development activities address the teaming issues?
> (2) Are there some types or categories of decisions or actions that the Human-Electronic Team should never be trusted with?
> (3) What oversight checks should be placed on the Team?
> (4) How does the Team communicate with the higher authorities?
> (5) Are there other issues besides teaming which are crucial to the operational application of the Electronic Crewmember concept?

The workshop comprised formal paper sessions and structured small group discussions. The proceedings are published as reports of the sponsoring laboratories.

# ORGANISING COMMITTEE

## CO-CHAIRS

John Reising, PhD.
WL/FIPA
Wright-Patterson AFB.
Dayton, Ohio 45433
United States

Robert Taylor
DRA Centre for Human Sciences
CHS 3
F 131, Rm 43
Farnborough, Hants, GU14 6SZ
United Kingdom.

## TECHNICAL COMMITTEE

George Brander, Maritime Command & Control Division, DRA Portsdown, UK.
Jeremy Clare, Cambridge Consultants Ltd., Cambridge, UK.
Terry Emerson, Joint Cockpit Office, USAF WPAFB, Dayton, Ohio, USA.
Howard Howells, Man Machine Integration Department, DRA Farnborough, UK.

## MEETING ADMINISTRATION

Jill Haugh, DRA Centre for Human Sciences, Farnborough, UK.
Joanne Rainey, DRA Centre for Human Sciences, Farnborough, UK.

## SPONSORS

## ACKNOWLEDGEMENTS

# WORKSHOP BACKGROUND

Ever since the movie *Star Wars* showed Luke Skywalker and R2D2 teaming up to destroy Death Star, there has been considerable speculation as to how an efficient pilot-robot team could be created. Since weight is a critical design factor in airborne systems, the literal building of a pilot-robot team has not been undertaken; rather, the emphasis has shifted to incorporating the intelligence of the robot. As work in this area progressed, such terms as "electronic crewmember" and "black box back seater" began to enter the vocabulary of both the crewstation design and computer software communities. While the use of these titles served to stimulate thinking in the area of human computer teamwork, a major program was required to build an electronic crewmember (EC); in the US this took the form of the Pilot's Associate (PA) Program. The establishment of the PA Program in 1985 gave credence to the idea that the building of the brain of R2D2, in some very simplified form, might be possible. Some of the results of this program have been transitioned to the US Army's Rotocraft Pilot's Associate Program which continues to strive for the same goal. In Europe, AI efforts have centred around a number of programs. These include the French "Co-pilote Electronique", the British Mission Management Aid (MMA), and the German CAMA and CASSY Cockpit Assistant Systems. They too have tried to achieve the goal of human computer teamwork in the cockpit.

In the next two years, numerous discussions were held to explore some of the cockpit ramifications created by the use of a pilot-EC team within the aircraft. These discussions occurred in various technical meetings within the US and Europe. In one of the meetings held in the US, attended by representatives of the Air Force of the then Federal Republic of Germany (FRG), as well as UK and US representatives, the idea of the initial workshop was born. Although progress on the idea of a workshop on human-computer teamwork

continued, in 1987 an event occurred which demonstrated the definite need for a workshop.

In April of 1987, USAF representatives gave a paper at a meeting of the Royal Aeronautical Society in London, and again at a meeting of the Ergonomics Society in Swansea, Wales. The subject of the paper was "Workload and Situational Awareness in Future Aircraft", and a section of the paper discussed workload sharing between the pilot and the EC. During both meetings the same kinds of questions were asked: Is the pilot always in charge? Can the pilot and EC really be called a team? Why do we need a pilot at all?

These thought provoking questions resulted in continued discussions with technical personnel in the US, UK and FRG, and the result was the 1988 workshop entitled, "The Human-Electronic Crew: Can They Work Together?" (RAF IAM BSD-DR-G4 Dec 88; WRDC-TR-89-7008). Following the 1988 workshop, interest was expressed in holding an additional meeting on the topic of human-electronic crew teamwork. The result was a 1990 workshop entitled, "The Human-Electronic Crew: Is the Team Maturing" (RAF IAM PD-DR-P5, April 1991; WL-TR-92-3078, July 1992). Both the 1988 and 1990 workshops were sponsored by the USAF European Office of Aerospace Research and Devlopment (EOARD), and hosted very generously by the German Air Force.

There was a four year hiatus between the second workshop and the present one. Events relating to the end of the Cold War caused a very dynamic environment, with many governmental reorganisations occurring on both sides of the Atlantic. After these events were sorted out, plans began to convene the third workshop. Once again, EOARD sponsorship was obtained, and as a result the present Workshop, which the Royal Air Force and DRA Centre for Human Sciences graciously agreed to host, became a reality.

The meeting was divided into two sections: formal presentations (papers) and workshop. The 27 papers covered five major categories: mission descriptions, knowledge engineering methodology, trust development, interface design, and systems integration. A summary of the ideas from the papers is given below.

## Papers

One of the key points made by the representatives of the aircrew community was that decision aids should *help* them make a decision by offering information related to their decision making criteria. The essence of the idea is that aircrew do not want to be tag along button pushers who are relegated to a secondary role. However, the missions are getting so complex and the crew sizes so small that some automated decision assistance, besides criteria presentation, is likely to be needed.

Another interesting comment came from the human factors community. Their impression is that the technology already exists to give us the displays, controls, and decision aids that are projected for current upgrades and future aircraft. The automation community replied that, while this may be true superficially, there were many problems still needing to be solved. Such issues as logistical supportability are key, for instance, to the selection of decision support methods. Other issues involve real-time and on-time processing.

## Workshop

After the presentation of the papers, the second half of the meeting consisted of a workshop. Its purpose was to form six teams to deal with AI technology and cockpit implications of the technology. The teams were composed of three technical disciplines represented at the conference - aircrew, crew station designers, and artificial intelligence experts. At the end of the workshop, each of the six team leaders presented the results of their deliberations. The details are documented in the workshop section of these proceedings; a summary is presented below.

There was a consistent message from the aircrew to keep them involved at the earliest stage possible. It is interesting that they were not against the incorporation of the EC into their cockpits so long as they received no surprises from the intelligent computer, and the aircrew was always in charge. The key point from the software designers was that many different portions of an EC exist but have not been integrated. Another issue raised is whether the integrated EC can run in real time. The human factors specialists were concerned with the specific means of building trust between the aircrew and the EC. Specific guidelines for successful teambuilding between the aircrew and the EC were given in a number of papers, and the consensus was that some real progress had been made in this area.

## CONCLUSION

The overall worth of the meeting can be summed up in the comments of one of the team leaders who stated that there is no question of the worth of the EC. The main issue now is to detail which of the components are mature and which deserve further effort. His team produced a preliminary attempt at a structure (Figure 2 in the Report of Group Discussions) which can serve as a framework.

Besides the technical information gathered, one of the major accomplishments was the positive interchange among the participants. There was a genuine sharing of information and ideas in order to attack the common problem of information overload in the cockpit. The participating countries are striving to reach a common goal, and the ideas exchanged in the meeting should prove beneficial to all of them.

# KEYNOTE ADDRESS

by

Group Captain G A Miller RAF
Operational Requirements (Air)
Ministry of Defence

I would like to thank the co-chairs, Bob Taylor and John Reising, for inviting me to give this keynote address to the International Workshop on Human-Computer Teamwork. Two aspects of the workshop are, to me, particularly significant. First, the international flavour will, I hope, provide a good cross-fertilisation of ideas. With the current scarcity of money and expertise, we must do all we can to improve cooperation in human-centered design as in other areas of research. Second, I think the broad spectrum of backgrounds represented here is essential. Without a dialogue between the research community, industry and the operators we could all too easily waste our resources searching for answers to altogether the wrong questions. Therefore, I see this is an important forum in which to share our views and develop ideas on the way ahead. As the MOD sponsor for human factors research, this is a process in which I am very pleased to be involved. Having been given the opportunity to open the debate, I will take a few minutes to outline my views on the significance of the teamwork approach to system design, and offer my thoughts on some of the many difficult issues that have to be resolved.

With the ever-increasing sophistication of avionic systems, we are fast approaching the situation in which the aeronautical industry can offer systems that are technically capable of meeting our operational requirements but which the aircrew cannot fully exploit; perhaps we are already there. Part of the problem is procedural. The procurement process brings with it a tendency to design avionic sub-systems before fully considering how they will be integrated into the cockpit and, because of the high cost of retrospective updates, we operators normally have to live with the result. The original Jaguar Navigation and Attack System is a good example of this weakness. The system was so poorly integrated that it is thought to have contributed to a number of fatal accidents, but it was almost 10 years before a more satisfactory upgrade was eventually installed.

In the UK, we have made some progress in this area, by putting human factors issues at the front end of the procurement process, and I know that other nations are making similar procedural changes. Unfortunately, there is a more fundamental problem that cannot be solved by the same method. The problem is that integration only ensures that elements of a system function **with** each other, which is not enough. We need to develop systems that will work **for** the aircrew to keep the pressures on them to a manageable level. In essence, future aircraft systems must **share** the intellectual workload, not just minimise their contribution to the overall task. As an operator myself, I can intuitively accept that this could be achieved by designing systems around a cooperative teamwork model but the concept raises many fundamental issues. Time prevents me from introducing all of them, so I will concentrate on the few which I feel are the more important.

Who should be the team leader - the mission computer or the human? What types of teams should the system emulate? How do we ensure that the team samples we experiment on are representative? What human characteristics should we allow for in our team? How many humans should there be in the aircraft? How much should the team members trust each other? And finally of course, can we trust the team? It is this last issue that we will be concentrating on this week but, of course, the issues are very much inter-related so I would like to provide a starting point for a wide ranging discussion.

Let me deal with the easy issue first - why aircrew must be in command. The argument is essentially the same as that for having a human in the cockpit at all. Modern computers,

particularly knowledge-based systems, offer a quite remarkable leap in the decision-making capability of the overall human-electronic system, but we cannot directly compare these machines with the human brain. We cannot even agree on a model of human intelligence yet, so we have no way of producing an electronic equivalent. (An alternative perhaps, but not an equivalent). Therefore, I believe it will be some considerable time, if indeed ever, before knowledge-based systems are capable of taking the intellectual lead in aircraft cockpits and, until then, I think we should be concentrating on the decision-support available from knowledge-based systems.

For those of you who are not convinced by this argument, there is another important area in which electronic systems are found wanting. Computers cannot replicate the ingenuity of the human mind that enables us to respond to unexpected situations in novel and unpredictable ways. I do not mean the ability to act in a random manner - my office computer seems to do that whenever I'm working to a tight deadline! I mean our ability spontaneously to grasp an opportunity, apply our imagination and creativity, and determine a course of action unique to the individual and the situation - a course of action that a computer might consider illogical. This is not a weakness but a great strength as it allows us to surprise the enemy, which is one of the key principles of war at all levels from grand strategy down to a brief air combat skirmish. Only when this unique characteristic is achieved in a computer do I feel that we should entertain any thoughts of replacing the human in the cockpit. In the meantime, I see the computer carrying out the more mundane task of flying the aircraft while the human, using his or her inherent insight, inference and intuition, fights the battle.

Hopefully that provides a starting position for any debate on the intellectual pre-emminence of the human. Let me move on to some of the other questions I raised.

Firstly, what type of team should the human-electronic crew emulate? An obvious starting point would be the crew of a multi-seat aircraft, but this can be a strictly heirarchical organisation working within a narrow set of teamwork principles. The crew structure has its roots in centuries of military tradition and allows officers to exercise command, control and leadership over their men. This may not be the best structure for a human-electronic crew so we may need to look elsewhere for a complete picture of the way in which our crew ought to operate. On the other hand, the more successful multi-seat crews often put aside the rigid structure of rank and position when they enter their aircraft and their overall performance benefits from a combination of the unique strengths of each individual. Operating in this manner, the performance of a crew will invariably exceed the sum of each individual's abilities.

We may have much to learn from non-heirarchical organisations, in which elements form ad-hoc groups to solve problems, each element being involved in different problems on a time-share basis. The overall aim of the mission is provided by higher authority, in this case by the Operational Commander. The crew leader then determines the best strategy to achieve this aim and directs the rest of the team accordingly but does not get embroiled in the detailed work of the groups. Of course, this approach brings its own issues, not least of which being the need to keep the aircrew informed in order to maintain situational awareness, the need for effective 2-way communication, the need for the behaviour of the machine to be predictable and, most importantly, the ultimate need to maintain the pilot's authority. These issues will be difficult to solve but it does illustrate the point that we need to cast the net wide in our research.

The follow-on question is about how to ensure that we use representative samples of teams during research. I note that one of the papers to be presented this week calls on an experiment using an experienced flight-crew in a simulator environment. This raises an interesting area for discussion. In my experience and as I mentioned earlier, the composition of a crew has a marked affect on the way it operates, from routine tasks such as inter-cockpit communication, right up to major tactical decision-making. Essentially, each member of the crew adapts his technique, and his degree of trust, to suit the experience and capabilities of

the other crew-member or members. Also, even a constituted crew (that is to say a crew who routinely and frequently operate together) can function very differently from one mission to the next depending on fatigue, distraction, motivation, and so on. Indeed, there were some genuine surprises amongst our constituted crews during the Gulf War due to the presence fear and in most cases it resulted in an even greater level of achievement. Therefore, later in the week you may wish to discuss how we will ensure that the teams used in our experiments are valid. Perhaps this discussion could be extended to address wider issues: should, perhaps, the electronic crew-member be optimised for a mythical standard aviator, or should we develop systems that can adapt to individuals? In either case, to what extent should the electronics cater for the vast array of human foibles?

Another major question is how many humans we will have in the aircraft. This topic could be the subject of a workshop in its own right so I will not attempt to cover it in any detail. However, I would like to make a couple of observations. Having flown operational tours on both single-seat and 2-seat aircraft, I have no doubt about the benefits of 2-man crews. I have found navigators to be invaluable for buying the beer and carrying your bags on detachment. Actually, there is a serious point here. The comradeship which develops in a constituted crew cannot be replaced by electronic systems. Again, the Gulf War reminded us that going into battle is a very daunting experience and the Tornado crews drew great strength from the fact that they had another human on board during the first few nights of action over Baghdad - the crews of the single-seat aircraft deserve even greater respect. There are problems with multi-crew operations, not least being communication, but there are also many operational benefits which can all be boiled down to the one over-riding point that the workload is **shared** in a multi-crew environment.

The debate on crew numbers has a direct relevance to this workshop. Firstly, human factors work has an important contribution to make in determining how many seats we should have in future aircraft. Secondly, I am quite sure that team dynamics varies with different

numbers of humans so our research may have to follow 2 separate paths at some stage. Finally, the number of humans will have a marked affect on the degree of trust that will be conferred on the team as a whole. This is the case now, and I see no reason for it to change significantly when electronics also form an integral part of the crew in the future.

This leads me on to the issue of trust. The first point is that we should not under-estimate the capabilities and needs of the human in the human-electronic crew. As a Tornado squadron commander, I was happy to authorise my crews to fly at low level through poor weather using the Terrain Following Radar because the pilot could monitor the system and over-ride it if necessary. I would not have trusted the electronics alone to conduct this critical task. More importantly, I did not need to - the aircrew were well trained with a highly developed survival instinct. In other words, my trust was earned by the **combination** of man and machine, not by the machine, or the man, alone.

On the other hand, if future electronic systems are to **share** the workload, we will have to reduce the pilot's involvement in the control loop for some tasks. This will require a detailed knowledge of the risks involved, coupled with some difficult decisions about the consequences that we are prepared to accept. One thing is certain and that is that we cannot expect to achieve **absolute** trust - neither the human nor the electronics will ever be totally infallible. As a general rule we operators will expect a very high level of confidence in potentially life-threatening situations, so we would expect to exploit the capabilities of both man and machine. Conversely, for routine tasks, in which the consequences of mistakes are tolerable, we will allow the machine to make decisions autonomously. Between these extremes, we may need to develop a range of predictable automation levels. We will also have to resist the temptation to err on the safe side as this will overload the pilot and prevent us from meeting our objective of cooperatively **sharing** the workload. This also raises questions about aircrew training in the context of when to trust and when not to trust the aircraft's automatic systems. Thirty years ago,

pilots got airborne clear in their minds that if it all went pear-shaped, they could trust their own judgement, were able to take over, and I believe that they genuinely expected the unexpected and were well-able to cope. Today's young men are, understandably, somewhat confused. Rather than expecting the unexpected and being prepared to take it on, many now simply fear that the unexpected will occur and are unsure how they will, or should, react. They have been taught to trust the automatic systems to the point in some cases that the automatics are more reliable than their own judgement. This reinforces the need for a dialogue between the operators and the research community if we are not to continue developing systems which leave the pilot guessing at a time when his thoughts should be clearly focussed on either fighting or problem solving.

The final point I would like to make is about the affect that increased trust will have on the way we operate. Currently, Commanders often feel unable to trust computers and other electronic systems sufficiently to allow them to act autonomously. As a result, restrictive rules of engagement are applied which can prevent systems being used to their full potential. For example, pilots may have to identify a target visually before attacking, to check that the sensors have correctly predicted that the target is hostile. However, the rules do not just reflect the Commander's confidence in his weapon systems. They are also the politician's final means of controlling events on the battlefield. Thus, even if we develop intelligent, predictable and trustworthy systems, this progress is more likely to translate into an increased confidence of success, rather than greater freedom of operation, and we cannot expect a sudden change in the way we go about our business.

In conclusion, I hope I have given you some food for thought, both over the next few days and afterwards, when you return to your normal work. I am very concious that I have provided more questions than answers, but I feel that this reflects the current situation quite accurately. For my part, I have an open mind about the issues raised by the human-electronic crew concept, apart from my views about keeping humans in an overall position of authority in the cockpit. I am looking forward to hearing your views on these issues and others that will arise this week. As a customer, I also look forward to reading the workshop report.

I very much regret that I'll be unable to stay for the whole week to join in the debate but I'm sure the other operators will be more than happy to offer their views. A word of caution though: it is an established fact that the number of opinions held by a group of aircrew is an exponential function, where 'x' is the number of wings on their chests. Thank you for your attention - I hope you have a rewarding week.

# SESSION I - MISSION SYSTEMS

## SYNOPSIS

The papers in this section describe HE-C systems and requirements for different missions. Operational roles and tasks covered in this section include the following: nap-of- the-earth attack helicopter (Paper 1); pre-mission planning and support for future tactical aircraft missions (Paper 2); target tracking in airborne maritime operations (Paper 3); planning, prediction, and target designation for precision guided munitions in battlefield air interdiction (Paper 4); dual pilot fixed-wing military transport aircraft (Paper 5); and high-level battlefield command, control and communication (Paper 6). *Paper 1* describes the latest techniques in "glass cockpit" design. The authors describe the design rationale for their attack helicopter cockpit, and they discuss the tools used to create and validate the system, from initial conception, through to fully manned mission simulation. Typical control display formats are described and illustrated. They report how human information processing principles of schemata and chunking have influenced the design of the cockpit management system. This paper demonstrates how prudent, mission-oriented application of automation can provide practical solutions to many current mission system problems, and sets the context in which applications of artificial intelligence technology will need to make a difference. The other papers argue that intelligent, knowledge-based systems have potential for technology to go further, and to provide assistance to the human operator, by aiding human decision-making, across a wide range of missions and tasks. Papers 2,3, and 4 describe the different kinds of assistance required for relatively specific, mission critical tasks involving planning, prediction and information management. *Paper 2* describes how future intelligent mission support systems will be required to be re-configurable in accordance with mission specific information. They will need to provide satisfaction of co-ordination requirements with other flight elements and co-operating intelligent systems, and they will need the ability to adapt to individual operator preferences for their intelligent system activities. *Paper 3* describes the requirement in target trackers for assistance in management of the sonics sensor environment to capture high quality target data, for assistance in the adoption or rejection of target data, for aid in the management of the information used by the tracker, and for assistance in the evaluation of tracker performance. *Paper 4* describes how tactical pilots need assistance in determining weapon delivery outcomes with stand-off weapons, including consideration of stringent collateral damage requirements, of pre-planned mission survivability data, and of cockpit imagery for manual target designation. *Paper 5* describes the nature of this intelligent assistance in an airborne system, and emphasises that the computer does not make autonomous decisions, and that it does not make weapons system interventions without orders from the pilot. *Paper 6* discusses inherent human decision-making limitations in the context of battlefield command systems, and anticipating the prospect of improving computer decision-making performance, it poses the controversial possibility of a totally automated battlefield. In general, at present the operational community want systems which provide assistance and advice in decision-making, but not second-guessing. They do not want EC to try to do better than, or to over-rule, the human decision-maker. They want EC to accept and to assist the decisions of the human operator; i.e. to help the operator make better (timely and tactically correct) decisions. Whether or not computers can out-perform human decision-making is a hypothetical question. At the present, and for the foreseeable future, when dealing with novel situations and decisions that are non-procedural, not planned, unpredictable and unexpected, it seems necessary to have human creativity, ingenuity, and flexibility in combat tactics to keep the winning edge.

# DEVELOPMENT AND EVALUATION OF THE AH-1W SUPERCOCKPIT

C.D. Holley
Principal Engineer, Crew System
Bell Helicopter Textron, Inc., Fort Worth, Texas


M.L.Busbridge
Consultant Engineer, Venom Management Team
GEC- Marconi Avionics Ltd., Rochester, England

## ABSTRACT

GEC-Marconi Avionics Ltd., and Bell Helicopter Textron, Inc. have proposed an extensively modified AH-1W SuperCobra for the British Army's new attack helicopter. Called Venom, the aircraft features an advanced technology mission equipment package (MEP) integrated at the human/machine level by the AH-1W SuperCockpitTM. This cockpit is one of the world's most capable and integrated attack helicopter crewstations, incorporating the latest techniques in "glass cockpit design." These techniques include liberal yet prudent mission oriented application of automation, enhanced data management that makes the right information available to the crew at the right time with improved geometric accommodation. The design rationale,in the creation of the SuperCockpit,to reduce crew workload and enhance mission effectivity at the same time as maintaining crew trust in the advanced mission orientated automation,are described herein.In addition the tools used to create and validate the SuperCockpit,from initial conception through to fully manned mission simulation, are also discussed;together with examples of typical control display formats.

## INTRODUCTION

GEC-Avionics (GEC) and Bell Helicopter Textron Inc. (Bell) have proposed an extensively modified AH-1W SuperCobra for the British Army's new attack helicopter. Called Venom, the aircraft features an advanced technology mission equipment package (MEP) integrated at the human-machine level by the AH-1W SuperCockpitTM. Jointly designed by Bell and GEC, this cockpit represents one of the world's most capable and integrated attack helicopter crewstations, incorporating the latest techniques in "glass cockpit design." These techniques include liberal yet prudent mission-oriented application of automation, enhanced data management that makes the right information available to the crew at the right time, improved geometric accommodation, and usability testing to reduce crew workload and enhance mission effectiveness.

The AH-1W Super-CockpitTM embodies a number of desirable physical attributes to facilitate mission success and safety. It has tandem crewstations that are configured in such a manner that duty functions are inter-changeable between fore and aft cockpits. While the preferred crew roles are pilot in front and copilot/gunner (CP/G) or commander in the rear, either crewmember can perform his duties from either crewstation. With some minor exceptions, controls and displays are functionally and physically identical in both cockpits, thereby improving logistics support, transfer of training, and mission effectiveness. Differences in between-cockpit arrangement of controls and displays have been minimized to the extent possible. Control display formats on the MFDs are straightforward and easily accessible. Flight control grips and mission grips are flight and mission oriented with hands-on access to critical functions. The cockpit geometry has been reworked for improved anthropometric accommodation for the targeted population. Crew vision has been emphasized, both out-the-window and sensor aided. External visibility is now considered superior to that of any other attack helicopter, with the SuperCockpitTM achieving approximately a 20% improvement over the current AH-1W. The Stability and Control Augmentation System (SCAS) provides proven aircraft handling qualities and will receive additional pilot-aiding functions. Mission and flight coordination between crewmembers is direct, simple, and positive with the SuperCockpitTM. Crew performance is improved through an overall design philosophy that encompasses augmented flight controls, pushbutton annunciators (PBA) for discrete control-display inputs, full alphanumeric keyboards with integral function keys and switches, and enhanced situational awareness via a digital map subsystem. A state of the art targeting system is displayed on an advanced technology display suite together with the use of wide FOV fully binocular helmet mounted displays. Standby flight instruments in both cockpits ensure that the crew is never without the necessary information for safely flying the aircraft should a major failure of both dual redundant mission systems occur.

Extensive cockpit automation and an exceptionally user-friendly interface give the crew the time required to attend to mission requirements rather than spending critical time dealing with system operations. For example, consider the situation where the pilot, flying nap-of-the-earth (NOE), spots a target of opportunity and designates it to

the CP/G. The CP/G, who has been communicating and replanning the route, presses the LOS ACQ switch (line-of-sight acquire) on his mission grip to slave the targeting system to the pilot's LOS, and with the target displayed on one of the two MFDs (TGT page) and with full "hands-on" control of the targeting system and weapons system engages the target.Crew coordination is efficient and effortless. Together, the crew collects and acts on information required to successfully complete the mission.

## SUPERCOCKPIT

### Mission Equipment Package
The UK attack helicopter is required to perform the following missions: anti-tank (primary), anti-personnel, air-to-air, ferry, instrument flight, training, reconnaissance, artillery observation, suppression of air defense, and limited search and rescue. Performing these missions includes the capability for day/night, adverse weather operations and nap-of-the-earth (NOE) tactics.

To support these requirements, the MEP includes: a head steered piloting FLIR system: a targeting system inclusive of TV/FLIR, laser ranger/designator, laser spot tracker, auto search and multiple target tracking; a data loader; an advanced aircraft survivability equipment (ASE) suite; and a fully integrated stores management system controlling a versatile weapons suite that includes a gun, air-to-air missiles, air-to-ground missiles, and rockets.

The primary control-display package is comprised of alphanumeric keyboards with collocated liquid crystal displays (LCD), mission and flight grips for hands-on functions, LCD MFDs, and helmet-mounted displays (HMD) with integral image intensification (I2). Touch screen and interactive voice technologies were excluded on the basis of cost and risk. The system is integrated together with two dual redundant mission computers,each machine driving a MFD in both cockpits, a HMD, a LFD and associated keyboard together with the capability of generating a digital moving map display from data derived from the mission loader cartridge.

### SuperCockpitTM Configuration
The SuperCockpitTM configuration is shown in 3-dimensional form in Fig. 1 and in two-dimensional layouts in Figs. 2 (forward cockpit) and 3 (aft cockpit). In addition each crewmember is equipped with an HMD inclusive of night vision intensiver capability with electro-mechanical head tracking. Each HMD thus being capable of steering either the piloting or targeting systems and displaying either sensor or the NVI imagary in conjunction with superimposed Heads-up symbology.This symbology is tailored to the flight regime and includes flight, weapons, systems, and CWA information.

### SuperCockpitTM Display Theory
From a human factors engineering (HFE) perspective, glass cockpits present a tremendous challenge for designing a system that deals effectively with managing the vast amount of information potentially available to the crew displays via the onboard computers and sensors. Having all these data available does not facilitate mission success and safety unless the crew has easy access to the correct information in a timely manner (Ref. 1). Either too much or too little data at the wrong time can be equally disastrous. Established models of human information processing (e.g., Refs. 2,3) were used throughout the SuperCockpitTM development to provide a "road map" for design decisions related to the human-machine interface. Two principles of human information processing were particularly relied on for assisting with the information management aspect of the SuperCockpitTM design and cockpit management system: schemata and chunking.

The concept of schemata originated with Kant in the 19th century and was introduced into psychology by Bartlett in 1932 (Ref. 4). Essentially a schema is an abstract, generic representation in human memory of an object, idea, process, or procedure. This abstraction contains slots or placekeepers that get filled (instantiated) when the schema is invoked (Refs. 2, 3). If an individual has developed an appropriate schema for a particular situation, then processing of the information associated with that situation can be facilitated by ensuring that instantiations coincide with preexisting slots. Standardization of formats between multifunction display (MFD) pages and access procedures for different MFDs (which are identical) represent two examples of the way the SuperCockpitTM design capitalizes on the schema approach.

The control logic for accessing the MFDs was also supplemented by the psychological principle of chunking. This concept was introduced by G.A. Miller in 1956 in his classic paper, "The Magical Number Seven, Plus or Minus Two" (Ref. 5). In essence, he demonstrated that short-term memory has a limited processing capacity that ranges from five to nine one-syllable words (the range also varies with differing stimuli in accordance with the limited capacity proposition). Miller further demonstrated that this processing limitation can be "overridden" by cognitive restructuring of information to make it compatible with the limited capacity. In other words the information can be organized into representative superordinate "chunks" that provide cues for retrieval of subordinate data by long-term memory. This is somewhat analogous to using an acronym as a mnemonic aid for

retrieving a word phrase. Miller's initial work has been verified and expanded by a number of researchers (Ref. 6). Other researchers have determined that failing to provide an obvious organization to the information creates inefficiency, as subjects spend unnecessary time trying to create one (e.g., Ref. 7). One example of chunking and organization applied to the SuperCockpitTM design was in canalizing MFD access via eight subsystems, further organized into two major subgroups (Fig 4).

One of the underlying philosophies associated with the SuperCockpitTM design was to avoid re-placing or redesigning an existing cockpit component unless such redesign was necessary to enhance mission effectiveness or safety, or to comply with system specifications or other customer requirements.

For example, the existing cockpit lighting control panels did not provide full functionality between cockpits and used magnetic toggle switches that experience had shown to be cumbersome in operation. For the SuperCockpitTM, the panels were redesigned to provide the same panel in both cockpits (there are minor sizing differences to accommodate installation requirements) and lighted PBAs are used for shared functions so that switches on both panels always indicate the correct lighting status.

Within the redesign constraints, maximum attention has been directed towards eliminating clutter in the crewstations. A primary goal in this regard has been to keep the instrument panels and all other vision blockages as small as possible, thereby maximizing external visibility and enhancing mission effectiveness and safety. Another goal was to implement panel controls and displays via MFD integration, as opposed to traditional dedicated panels/switches.

**SuperCockpitTM Equipmements**
1. Two high-resolution, color LCD MFDs are mounted side by side directly in front of each crewmember. They are approximately 27 inches forward of the design eye point (DEP) and symmetrical around the center line. The MFDs have a 8 x 6 inch display surface surrounded by a bezel that contains 26 switches for control-display interface locatedwithin Zone-1 reach. Eight of these are dedicated,engraved,"hard-key"switches, organized in two groups of four along the bottom of the display. Each of these switches corresponds to a control-display subsystem as identified by its legend and when selected provides all the pertinant display and control for that sub-system. This is provided by "hands-on" controls (cyclic and collective for the pilot functions,and mission grips for the CP/G) together with the remaining 18 pushbutton switches which are

"soft-keys" orline-addressable control switches wherein their function and label is specific to each display page. All normal fly and fight functions performed routinely by both operators can be observed and controlled from the top eight display pages. Individual display sub-pages accessable from the top pages being used to set up the various sub-systems in the event of the mission loader being unavailable, a change to the pre-planned data cartridge loaded mission, a more detailed display required, or for maintainance crew use. The normal display controls of brightness, contrast etc.are located at each corner of the bezel.

Figures 5 to 9 inclusive are examples of 5 of the 8 top pages currently designed and presently being evaluated on the respective BHTI and GEC simulators.

2. The limited function display (LFD) is an LCD that is collocated with the keyboard. It is used to display four types of data: keyboard input, CWA messages, inter-cockpit status, and subsystem status. Information display is organized as follows:

   a. The bottom line serves as scratch pad for the keyboard.

   b. The next 3 lines are used as part of the CWA alerting system.

   c. The next line provides inter-cockpit status, toggling between weapons and communications subsystems.

   d. The remaining 75% of the display is used to toggle between a Remote Frequency Display (RFD) for communications status and a pictorial display of armament status.

3. The keyboard selected for the SuperCockpitTM has extensive and favorable military experience onboard the OH-58D. While collocated, it is not integral with the LFD. In addition to a full set of alphanumeric keys, it also contains three toggle switches and five pushbutton switches for implementing selected dedicated functions such as CWA and emergency communications interactions.

4. The forward cockpit's cyclic control stick is located on the right console. The cyclic grip contains nine switches to provide hands-on control of the following functions: radio/intercom transmit, force trim, weapons select, SCAS disengage, HMD video (FLIR/I2), weapons action/steer,missile cage/uncage, display select and weapons fire (trigger). The shape of the cyclic grip and the location of switches is currently being defined using man-in-the-loop (MIL) simulation. The

aft cockpit uses the same grip mounted on a kneeling, center-stick. The kneeled position provides clearance for using the mission grips in their non-stowed position.

5. The collective control stick is located on the left side of the crew station and contains two twist-grip throttles for engine power management. A conformal (shaped in accordance with human engineering considerations) collective grip is located on the end of the collective stick. This grip contains eight switches to provide the operator with hands-on control of the following functions: radio frequency select, radio select, idle stop release, emergency jettison, countermeasures, search light control, search light slew and hover hold. As with the cyclic grip, the shape of the collective grip and the location of switches is currently being defined using MIL simulation.

6. Each cockpit contains two mission grips, located below the MFDs, that are installed on telescoping platforms. The grips pivot and rotate to an upright orientation when moved from the stowed to the operational position. In addition, the telescoping mount provides five lock-type positions for accommodating fore and aft adjustment. The left mission grip has nine switches for controlling the following functions: TV/FLIR focus, TV/FLIR gain and level control, laser fire (trigger), LOS acquire, FLIR polarity, track box size adjust,sensor select,action steer and FOV select. The right mission grip has nine switches for controlling the following functions: weapons fire (trigger), weapons select, turret/cursor slew, track function select,gun targeting select, FLIR Auto initiate/manual,HMD video, missile cage/uncage and weapons action/steer.

7. The forward cockpit contains four integrally illuminated standby flight instruments that, with the exception of being powered via the battery bus, function independently from all other aircraft subsystems and sensors. These instruments are barometric altimeter, airspeed indicator, attitude indicator, and magnetic compass. The first three of these instruments are duplicated in the rear crewstation; the instruments are located on the right side of the instrument panel in each cockpit and are easily viewable from the design eye position. The magnetic compass is located on the left side canopy rail in the forward cockpit and is viewable by either crewmember.

Other features of the SuperCockpitTM design include dual rearview mirrors in each crewstation, instrument panels moved closer to the design eye position, and an improved SCAS that includes attitude, altitude, and hover hold modes. The rear cockpit houses the data transfer module (DTM). In addition, geometric accommodation of the target population has been significantly increased by incorporating a four-way adjustable seat in the aft crewstation. The forward cockpit takes advantage of a recent modification to the canopy that, among other things, improves head clearance.

## CONCLUDING REMARKS

When fielded, the SuperCockpitTM will likely represent the most advanced attack helicopter cockpit in production.The development of the design in the timescales acheived would not have been possible without the capabilty of the interactive use of the GEC/BHTI simulators with their respective rapid prototyping capabilities.

A rapid prototyping capability is such an important tool that the design of a glass cockpit should not be undertaken without one. (This also applies to the capability for performing high fidelity MIL simulation.) In addition to the direct facilitation of the work of the design team, the rapid prototyping tool also serves as a device for improving communication and documentation outside the team. For example, control-display formats were down loaded as graphics files from the SG workstations and directly imported into PC-based desktop publishing software for producing program documentation. This communications capability is particularly beneficial when all members of the design team are not colocated.

The authors also acknowledge the major contribution made by the many USA and UK service personel who have "flown" in the various configurations of the respective simulators whose observations and feedback have validated the SuperCockpit design.

SuperCockpitTM is a trademark of Bell Helicopter Textron, Inc.

## REFERENCES

1. Graf, V.A., and Holley, C.D., "Human Factors Impact on the V-22 Osprey Cockpit: An Overview," American Helicopter Society, 44th Annual Forum, Washington, D.C., Jun 1988.

2. Holley, C.D., and Dansereau, D.F., "Networking: The Technique and the Empirical Evidence," in Spatial Learning Strategies: Techniques, Applications, and Related Issues, C.D. Holley and D.F. Dansereau (eds.), Academic Press, New York, NY, 1984.

3. Rumelhart, D.E., and Ortony, A., "The Representation of Knowledge in Memory," in Schooling and the Acquisition of Knowledge, R.D. Anderson, R.J. Spiro, and W.E. Montague (eds.), Erlbaum, Hillsdale, NJ, 1977.

4. Simon, H.A., "Information Processing Models of Cognition," Annual Review of Psychology, 30, 1979.

5. Miller, G.A., "The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information," Psychological Review, 83, 1956.

6. Baddely, A.D., The Psychology of Memory, Basic Books, New York, NY, 1976.

7. Tulving, E., and Donaldson, W. (eds.), Organization of Memory, Academic Press, New York, NY, 1972.

Figure 1

10

Figure 2  Two-dimensional Layout of Forward Cockpit

Figure 3  Two-dimensional Layout of Rear Cockpit

Figure 4 MFD Bezel



Figure 5 Display page SYS

12

Figure 6 Display page FLT
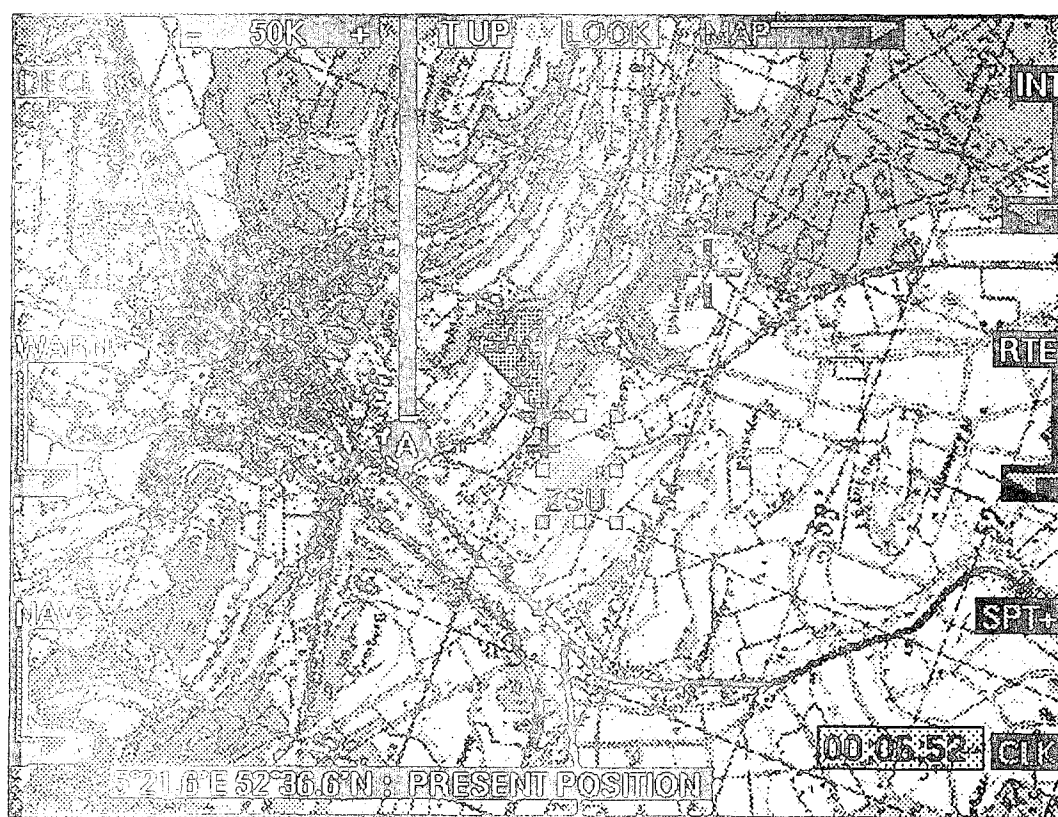


Figure 7 Display page MAP

Figure 8 Display page TGT
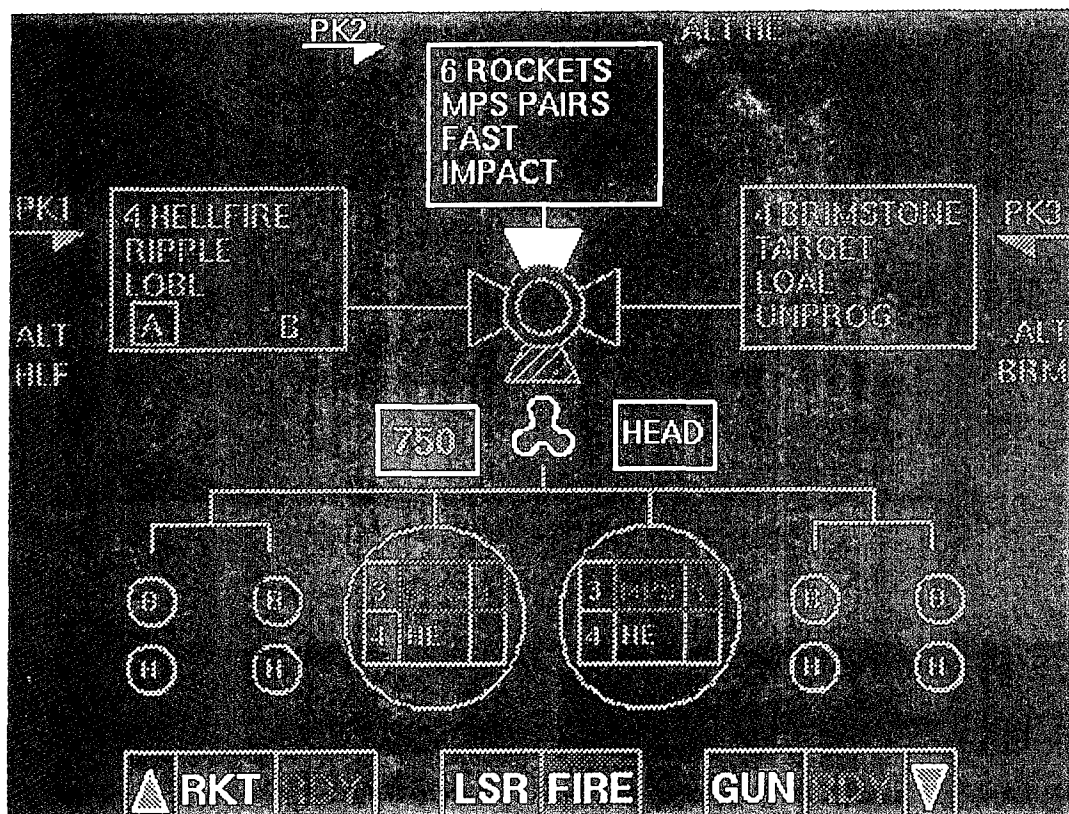


Figure 9 Display page WPN

# INTELLIGENT SYSTEM OPERATIONAL SUPPORT REQUIREMENTS

Thomas D. Aldern
GreyStone Technology, Inc.
15010 Avenue of Science, Suite 200
San Diego, California  92128

## SUMMARY

Next generation mission support systems are being designed to satisfy such conventional premission planning requirements as route planning, threat and countermeasures analysis, tactics planning, visual familiarization, combat folder preparation and aircraft computer and weapons initialization support. As intelligent systems technology matures and is embedded in current or next generation aircraft, additional premission operational support requirements will have to be satisfied. These include providing aircrew the ability to configure an intelligent system in accordance with mission-specific information, defining coordination requirements with other flight elements and cooperating intelligent systems, and establishing individual preferences for their own intelligent system activities. These capabilities will be critical to aircrew acceptance of their electronic crewmember as it allows the intelligent system and other aircraft avionics systems to be "tuned" to individual aircrew implementation of specific mission requirements and employment doctrine. This paper describes a premission planning system that supports these unique requirements as demonstrated in the Lockheed Pilot's Associate program.

## INTRODUCTION

Early success in development of the Pilot's Associate (PA) Program's complex knowledge-based, pilot aiding system has provided the opportunity to integrate near-term embedded avionics applications of this "associate" systems technology. The PA (figure 1) has developed as a set of cooperating, knowledge-based subsystems: two planner, two assessor, and one Pilot Vehicle Interface (PVI) subsystem. The program has served as a framework for demonstrating what associate systems technology requires in the way of new approaches for design, development, implementation and evaluation of cognitive-like functionality in avionics. This has led researchers in human factors to focus on the issues of pilot acceptance, trust, and human-electronic crewmember teamwork. As the associate system concept took form and pilots began to examine the more critical aspects of PA operation, the question of associate system control versus pilot-in-command became a key issue. Thus, developers adapted a pilot-centered operational philosophy which asserted "The pilot is in charge" and "The effort required to command the PA must be less than the effort saved by the PA" (Lockheed 1990).

This set the stage for debating pilot acceptance of the PA's aiding functionality because "the value of an aid is likely to be judged relative to unaided performance" (Rouse 1988), is not enough for the pilot. Merely implementing functionality within the framework of this operational philosophy lacks a certain sufficiency from the pilot's perspective. The pilot still needed the mechanism to assert

his command over the PA.
Previous research in this area
(Morris & Rouse 1986) concluded
that task performance can be
significantly improved if users
are in charge of their decision
aids. Separate research (Rouse
1988) subsequently indicated
"perceived ease of use of an aid
is also affected by the multi-
task nature of complex systems.
At the points in time when an aid
is most needed, it is likely that
pilots will have few resources to
devote to interacting with the
aid." In other words, if pilots
had the resources to direct and
monitor the associate system
during a critical mission task,
they would probably not need the
associate system to perform that
task anyway. This challenge is
pertinent to adaptive aiding
functionality in the PA PVI as
well as adapting or tailoring the
PA off-line and prior to the
mission.

Pilots perceive and react to the
Pilot's Associate system
differently from other "avionic
systems", and attribute the human
characteristics of skill and
knowledge to its performance
(Smith 1990). It is intuitively
obvious, that if the PA and the
pilot are to be successful they
must form a team. Early in the
PA program Reising (1985)
described the ideal team as
having, "such intimate knowledge
of how to work with each other
that they function as smoothly as
an Olympic figure skating pair,
each anticipating the moves of
the other while striving for the
same goal". Implications for
associate systems which provide
embedded skills and knowledge are
for the inclusion of friendly
human attributes to support pilot
acceptance of the system.
Therefore, to ensure mission
effectiveness and well
orchestrated teamwork between the
pilot and the associate system,
the pilots need a mechanism to
communicate team goals to the PA.

Domain experts recognized the
importance of the detailed
briefings that occur between a
pilot and backseater and this
evolved into the Mission Support
Tool (MST) subsystem of the
Pilot's Associate. This paper
discusses the MST development and
implementation in Phase II of the
PA Program. It emphasizes the
pilot's ability to partition,
allocate, and authorize tasks
which tailor the associate
system's functionality to improve
situation awareness and mission
effectiveness.

## THESIS

The Mission Support Tool is a
mechanism necessary to
successfully team and interface
the human pilot with the
electronic crewmember (PA). The
MST (figure 2) is a ground-based,
pre-mission interface between the
fighter pilot, a larger global
mission planning system, and the
associate system. As depicted in
figure 3, the primary function of
the MST is to furnish pilots with
a means for tailoring the PA to
their individual preferences for
air-to-air combat missions. The
MST keeps the pilot in command,
which domain experts agree will
pay high dividends in pilot
acceptance, trust, and human-
electronic crewmember teamwork.

Pilot acceptance and trust is
achieved only through the
repeated use of that aircraft and
weapon system element in as many
of the demanding circumstances in
the operating environment as
possible. In the case of
conventional avionics systems,
that means direct, hands-on
evaluation of the utility of the
system by the pilot. In the case
of another crewmember that shares
some of the mission workload, it
means the continuous development
of personal work relationships
with that person. In the case of
the associate system and other

16

intelligent systems, it means a combination of both approaches.

In developing pilot acceptance and trust regarding the associate's capabilities as a "supporting" crewmember, pilots need to be in command - physically and psychologically. A pilot can never afford to be surprised by what any other crewmember, including an "intelligent" associate system, will do. Further, the associate system must adapt to the pilot's personality for optimum crew coordination and maximum tactical effectiveness. And that means adapting to each individual pilot. Pilots need to know that associate systems will adapt to their own preferences, not the other way around.

## SOURCES OF INFORMATION

The MST achieves three objectives: (1) It provides the pilot authorized access to all controllable actions and functions of the PA subsystems for the purpose of setting desired authority limits. (2) It permits the pilot to develop or modify plans used by the PA during the mission. (3) It correlates and integrates information from critical sources and presents them to the pilot in intuitive form. The MST does not interfere with the PA inference mechanisms, and the PA will operate without it. However, just as a human crewmember's effectiveness will be enhanced by briefing him prior to flight, so will PA operation be more "skillful" with MST providing tailoring and specialization of the associate system prior to its use on any specific mission.

The MST subsystem overview (figure 2) shows the primary inputs to and outputs from the MST. The inputs are used to initialize the subsystem for a pilot planning/briefing session.

At the end of the session the selected outputs are used during PA initialization. A user interface allows the pilot to perform engagement planning and preference selection functions, including specialization of engagement plans, setting plan authority levels, reviewing mission plans, simulation of specialized plans and aircraft data transfer control.

## FINDINGS

Future avionics that employ intelligent or associate systems technology will have major impact on the human interface aspects of a total aircraft weapon system, whether it has a direct man-machine interface or not. Human factors will become an important part of the design of components and subsystems that traditionally have not required them because of no direct man-machine interface. These systems will require the inclusion of "friendly" human attributes and other human factors to support pilot acceptance of the system.

It is essential that human standards be considered in associate system design iterations to meet a different set of acceptability criteria - traditional system performance criteria and specifications are insufficient. Operational and tactical considerations have a more direct and continuing influence on associate system designs. With the PA for example, it is as if a "nugget" pilot is being trained and molded as he proceeds through the training pipeline on his way to earning his wings. In the PA program, developers established an operational task force that meets regularly with the systems and design engineers to review PA progress from the pilot's perspective, with notable success.

17

Pilots need to be in command of an associate system, therefore a means is provided for them to "tell" their intelligent system(s) what they want done, how, when, and to what extent it is to operate "independently". Associate systems must adapt to the personality and preferences of the pilot, therefore, a mechanism must be included with the system to allow this.

Associate system adaptability has been addressed in the PA program partially through the evolution of a Mission Support Tool that allows the pilot to "reach into" the various subsystems of the PA and tailor various aspects of its operational software to provide for:
(1) Tailored engagement plans.
(2) Specialized avionics and weapons employment plans.
(3) Extensive authorization levels, actions and priorities.
This planning collectively culminates in the establishment of a "personality" that matches the preferences of individual pilots.

The concept is analogous to a pilot conducting a face-to-face pre-mission briefing with a backseat crewmember before flight. All coordination is discussed and "contracts" made, workload assignments delegated, potential or preferred engagement options are identified, and priorities established. A McDonnell Aircraft Company study (1989) of operational pilots, indicated that there is a large variation between the views of experienced and less experienced pilots regarding use of such cockpit decision aids. Although not part of the original PA program, the MST is now an integral part of the approach to dealing with the "intelligent" aspects of the associate system and the pilot's need to have the system adapt to him or her.

A major realization within the PA program is the impact of intelligent systems "behavior" and other human attributes that will be embedded in aircraft avionics systems with associate systems technology. This will require an increased and continuing role for human factors engineers and operationally experienced personnel beginning very early in the design phase of those systems.

DISCUSSION

The MST Phase II prototype has addressed the preliminary issues of configuring intelligent mission software (PA) through a ground based, pre-mission specialization tool, therefore providing pilots with system predictability. Although, a rapid prototyping methodology and object-oriented approach were key factors contributing to its initial success, the MST's true payoff - pilot acceptance, trust and teamwork - was achieved during the PA manned system evaluation.

Smith (1990) points out that, "the only thing that can be said with certainty about the PA is that if the PA makes its home in a military airplane, then someday both the pilot and the PA will be in combat together". And further, "the challenge to builders of intelligent avionics systems is to ensure that the pilot and his PA-like system are not in combat with each other, but rather are on the same side and against the common adversary outside of the cockpit".

It appears that the challenge in achieving pilot acceptance of PA-like systems may be in the recognition of the "cultural" differences between the human, cognitive-like functionality of the system and pilots with varying experience. The MST is a tool that effectively deals with

18

these individual differences and
preferences of pilots.



Figure 1 - Pilot Associate Overview
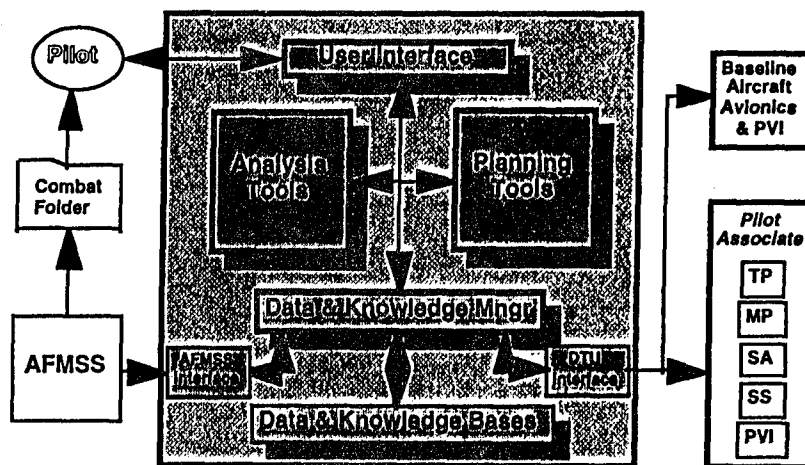


Figure 2 - Mission Support Tool

**Associate and Advanced Avionics Systems Should Support Mission and Personal Preference Specialization**

Operational Flight Program and Associate System Software Should be "Missionized" to Force-wide Tasking and Planning Data

Associate System Software Should be "Electronically Briefed" on Flight Element Contracts and Mission Specific Tactical Gameplans

Associate System Aiding Functions Should be Allocated, Partitioned and Authorized by the Pilot

Associate System Software Should be Robust Enough to Function Without Pre-mission Specialization

**The Pilot Controls the Length and Depth of a MST Planning Session - Not the Other Way Around**

The Pilot Can Manage and Interact With Single or Multiple Planning Functions Driven From a Single Top-Level Menu

The MST Will Automatically Load Pre-defined, Individual Pilot Default Preferences

An Pre-defined DTU File Load Can be Accomplished by Two Pilot Actions or Quickly Modified

Current AFMSS Mission Data is Automatically Loaded During MST Start-up

The Pilot Is (Still) In Charge!

Figure 3 - MST Concept of Operations

# Assistance to the Human Management of Target Trackers in Airborne Maritime Operations

I. S. MacLeod
Aerosystems International
West Hendford
Yeovil,
Somerset BA20 2AL UK

## SUMMARY

The human management of Target Trackers (TT), as applied in airborne maritime operations, is onerous and incurs a high operator workload, often at the most inopportune times. The majority of TT management tasks are highly structured with explicit rules making them suitable for assistance through the use of KBS techniques. The paper argues that this management area is highly suited to the application of Knowledge Based Systems (KBS) to assist the human.

## Introduction

Airborne maritime operations consist of military operations over the sea that can encompass search for targets, target tracking and their attack. Target Trackers (TT) are engineered aids to assist maritime tacticians in the accurate tracking of targets. TT have been in use in airborne maritime operations for at least forty years and their engineering design has become more and more sophisticated during that period.

The burgeoning sophistication of TT is usually argued as necessary to equate improved target performance capabilities and to cater for the greater accuracy and data rates of the sensors used to detect and follow targets. In addition, it is argued that sophisticated TT are required to allow the accurate and timely fusion of data from diverse sensors into target tracking information and the amelioration of tactical uncertainties caused by target deception strategies such as manoeuvre and the use of decoys.

These arguments will not be disputed by this paper. What will be disputed is the amount of assistance that the engineered sophistication of TT has provided to the tactician considering the operating management overheads that they incur.

The use of KBS will be discussed as a means of decreasing the tactician's TT management overheads, largely incurred by the need to search many lengthy lists and cross-refer the findings - the tactician's skills are best developed and employed elsewhere.

The intention is not to suggest means of directly assisting the tactician in the performance of tactical decisions or in the tactical control of the aircraft; if such assistance could be easily and meaningfully achieved, the tactician's cognitive processes would have already been successfully mirrored in engineered system design. The intention is to suggest means of assisting or relieving the tactician of onerous but structured tasks, tasks that can place high cognitive workloads on the tactician at the most inappropriate times and thus adversely effect their tactical appreciation and situational awareness (as the author can confirm). However, the properties of these tasks will support the application of a KBS approach to provide assistance to the tactician.

The approach suggested here is a requirements driven approach suggesting that KBS application might fulfil the requirement. It is not technology driven where the suggestion is that KBS technology is a panacea for all ills (discussion on the difference see Ref 1).

The subject is the management of TT. Such management encompasses the use and management of sensors, and sensor data, as well as the management of the actual TT in use.

## Sensor Management

In maritime operations, the simultaneous employment of sensors to cater for surface and sub surface targets has always been a problem because of the difference in the optimum aircraft heights for the use of each type of sensor and the difference in the form and quality of data obtained from each sensor. This is a problem in the tactical employment of the aircraft that is acknowledged but will not be discussed further in this paper.

This paper will concentrate on TT management and data fusion problems considering the primary use of one form of sensor; in this case the use of sonobuoys and a sonics sensor to track underwater targets by sound. However, the basic problems associated with the use of sonics are highly similar to the problems associated with the use of other sensors e.g., use of Dipping Sonar or Radar whether singly or in combination.

## Management of Sonics

Management of data derived through sonics equipment is essential to ensure that suitable data is produced for use by the TT. Unfortunately, this management frequently has to be performed

at the same time as the management of the TT and whilst the tactician is under stress e.g., time pressures and rapid performance of decisions associated with the tactical direction and control of the air platform and the deployment of aircraft stores and weapons.

To track a target with sonics requires that the appropriate types of sonobuoys are accurately placed in the water in order to glean data on the target. Target related data is then used to determine target identity, a line of bearing on the target, a fix on the target or data related to target performance. This sonics data is then filtered and fed to the TT which then converts that data to tracking information in the form of target course and speed, depth and indications of rates of change / accuracy of that information.
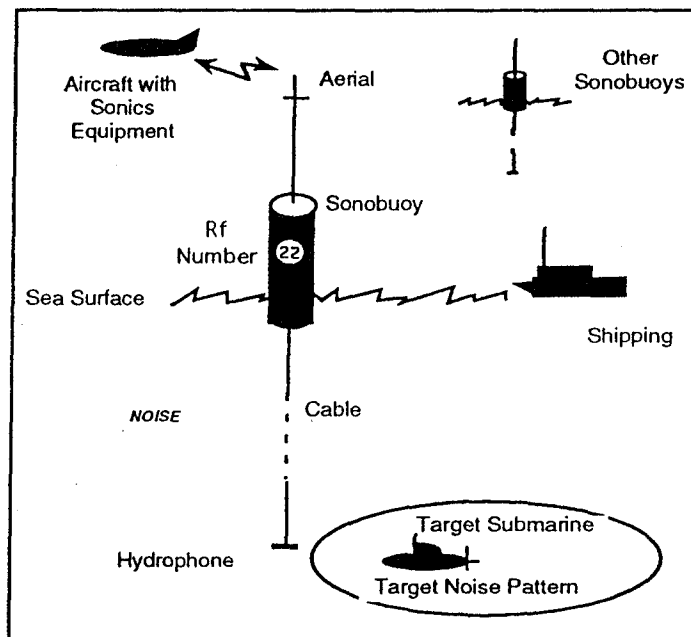


*Figure One - A Basic Maritime Scenario*

The operator must continually assess the accuracy of TT performance against cognitive derived but skill mediated assessments on sonics evidence. This assessment is necessary as tracker derived information is not only used to indicate the probable present position of a target, it is also used in the estimation of future target position in order that weapons and stores may be optimally employed and target tracking maintained. If the tracker has an accuracy bias, the operator must correct for it in their tactical appreciation of target movement and intentions.

There are many tactical implications that have to be considered in the use of sonics e.g. detection ranges on the target, appreciation of meaning and quality of target data, oceanographic conditions, target speed etc. However, there are in addition four main problem areas of management associated with the sonics tracking of an underwater target as aided by a TT, namely:

    i.    The management problem of correctly selecting, preparing for drop and dropping of sonobuoy(s) - Sonobuoy Stores Management.

    ii.    The management of information obtained from the sonobuoy(s).

    iii.    The management of the information utilised by the tracker.

    iv.    The assessment of tracker performance.

Each of these problems areas will be considered in turn.

## Sonics - Problem Areas in Management of Target Trackers

### 1 Sonobuoy Stores Management

A sonobuoy is a store that is positioned on the surface of the sea as a means of obtaining and relaying data on targets to a surface or airborne platform. Buoys can be used for the acoustical detection of sub surface and surface targets (both actively and passively), determining the sea ambient noise and temperature structure and for communication between units above and below the sea surface. Regardless of the type of buoy, the relay of data is achieved through the use of a set of radio or RF channels.

The problem with the management of sonobuoys resides in the many different parameters that have to be considered by the tactician prior to any sonobuoy deployment.

These parameters include (Ref 2):

a.      Currently almost 100 RF channels in use;

b.      Over 10 types of sonobuoys in current use in the 'western world', with the possible type utilisation per aircraft sortie reaching up to seven;

c.      A variety of settings for each sonobuoy depending on type and intended usage (e.g. life, hydrophone depth);

d.      Aircraft sonobuoy load variations (total numbers, types, RF channels, launcher capabilities);

e.      Sonobuoys in water / sonobuoys in aircraft?;

f.      Restrictions on the use of certain channels;

g.      RF channel Artefact and Conflicts;

h.      The mixture of buoys required to perfom a specific tactic;

i.      Plans and changes to plan;

j.      Buoys loaded in launchers / buoys required in launchers.

All the above parameters have to be considered prior to any sonobuoy drop and with relation to the target tracking requirements and the expected sonobuoy performance. Current sonobuoy stores management aids are based on lists; lists of stores in water, lists of stores in aircraft, lists of stores in buoy launchers. Usually, the lists give some indication of any inter list conflicts but none fully consider all the parameters listed above. Moreover, lists are difficult to search if the tactician is under time pressure to complete a set of tasks.

The consideration of the sonobuoy parameters listed above involves the tactician in unwanted 'mental arithmetic' at a knowledge based level (Ref 3). The result of the above is that sonobuoy management is often only easily and effectively performed pre mission, with changes required by actual mission performance being managed with difficulty. Frequently, 'home made' china graph boards are used to assist the tactician in this management process.

## 2      The Management of Sonics Information

The data obtained from sonobuoys is a mixture of good and bad target associated signals, signals emanating from other contacts both sub surface and surface, noise (e.g. produced by marine life, seismic disturbances, surface weather) and machine artefacts. The main role of the sonics operator is to filter out the identified target related data from the rest. To assist the sonics operator and the tactician, machine related assessments of the quality and accuracy of the data are also used.

Unfortunately, a second filter on the sonics data is required as:

•      Some sonics data is hard to associate with a particular target and requires corroboration from other sonics data or from its association with target information obtained by other means;

•      Some sonics data may appear to be of high quality but may have an inherent bias degrading its accuracy;

•      Machine related assessments of the quality and accuracy of sonics data are variable because of inherent difficulties in the engineering specification of the strong but variable influences of target and environmental based effects.

The above second filter is provided by the maritime tactician and is not effectively aided by current tactical systems.

## 3      The Management of Information Utilised by the Tracker

All engineered TT are limited in the data that they can handle at any one time. This limitation is usually coped with by replacing the oldest information with the newest. However, frequently the most accurate and pertinent data will have update rates and biases that are different from other data in use. To prevent the loss or degradation of best data requires careful management of what data is in use by the tracker.

Further, as sonobuoys are deployed, new data from these sonobuoys has to be appreciated by the tactician prior to managing the incorporation of that data into the tracker, frequently when there is little time available for that management or for tactical appreciation.

The assessment of best data is as a result of tactical appreciation of the available evidence. However, to force the tracker to appreciate that best evidence is usually associated with a high management workload by the tactician. The tactician is continually having to remove and add data obtained from particular sonobuoys as the target manoeuvres and target related sonobuoy data changes in form and quality. He is given little assistance by the engineered system.

## 4      The Assessment of Tracker Performance

It has been argued that the management of sonobuoy derived data by the tactician, in the assistance of tracker performance, incurs high management overheads. What has then to be considered is whether the payoff of these overheads is worthwhile.

Unfortunately, the drive to produce sophisticated TT in airborne maritime tactical systems has approached the target tracking problem from a predominately engineering standpoint assuming that a greater tracking accuracy will result and that this will be of obvious benefit to the human operator. Whilst the increased sophistication of TT has been manifest in their ability to use diverse

forms of sensor derived data, the accuracy and quality of that data is assessed by machine algorithms and takes little account of any characteristics associated with the operating environment and the particular target.

Furthermore, the high workload overheads involved in the tactician's management of target related data, data needed by the TT to efficiently operate, must promote the probability of data being mishandled and the tracker operating well below its optimum performance. It must be emphasised that the TT is primarily intended as a tool to assist the tactical appreciation of the tactician.

If the management of that tool involves a large proportion of the tacticians available time, the tactician may end up 'fighting the machine' rather than the enemy. Add to a degraded tracker performance the unbounded effects from real world uncertainties, and the result is that the human operator's trust in the TT is not fostered.

In particular, the following will be readily apparent to the operator:

- The limitations of the tracker's handling and following of target manoeuvre at the expense of the tracker's benefits;

- Overall tracker performance indices failing to live up to operator's expectations when compared to the tracker performance in reality:

- The high workload incurred in the tracker associated management for a perceived low value return;

- An obvious detraction in the operators appreciation of the tactical environment and situation seen as caused by the workload overheads incurred by tracker associated management.

To foster the tactician's trust in the tracker performance, it must be possible to tune and improve the tracker with relation to the performance it achieves in reality. It must be possible to determine and ameliorate tracker bias.

This determination of tracker bias, and its management, can be approached in two ways: 1) through a study of overall tracker performance through analysis of trials results and through simulated assessment of tracker performance using data derived from actual operations; 2) through allowing the tactician to tune the target tracker during actual operations to remove the majority of apparent tracker bias. This latter bias determined by considering tracker derived information on the target (e.g. target position, course and speed) with target related information achieved by other means (e.g. aural determined close pass on a sonobuoy, high quality manual fixing by doppler, course and speed as determined by the penetration of several sonobuoy barriers, mast sighting).

## What is Really Wrong with Methods of Current Tracker Management?

The following is a list of some of the problems associated with the methods used in the management of sonics data for current target trackers. Unlike the extensive lists associated with the management of TT in airborne systems, it can be perused at the reader's leisure.

1) TT design has not adequately considered the usability of TT by the human.

2) Onerous and structured management tasks, associated with TT, place high cognitive workloads on the tactician.

3) The burgeoning sophistication and complexity of TT has been associated with increasing cognitive loads on the operating tactician.

4) The increased sophistication of the engineered design of the TT has not been accompanied by a parallel improvement to its efficacy in practice.

5) Currently, TT is poor in its handling of noisy data and uncertainty. This is partly due to the engineered solution ignoring the sources of the problem in reality.

## Suggested Solutions to Ameliorate Management Overheads

The following are suggested solutions to ameliorate the management overheads in the use of TT.

### Buoy Load Management Aid

An KBS could be constructed to assist the tactician in buoy load management as the rules on the settings and use of sonobuoy types, both singly and in combination, are well defined. However, for the human load management involves many repeated and time consuming simple calculations, sorts and comparisons of results.

The KBS would have to utilise the following:

a. Knowledge of buoy load (e.g. bathythermal, sonobuoy etc.), any type conflicts and the details of the aircraft buoy load;

b. Knowledge of mission requirements / restrictions on buoy usage;

c. Knowledge of buoy types / numbers required for specific tactics;

d. Knowledge of current and future sonobuoy environment and associated sonobuoy settings;

e. Preload of rules / buoy usage plan;

f. Knowledge of current and required launcher load;

g. Ability to easily change plans and rules of aid;

h. An up to date knowledge of RF channels occupancy;

i. Priorities in buoy usage and type allocation to tactics;

k. Simple inputs indicating the tactician's intentions (e.g. Tactics A, B, C etc.).

The system would be designed to provide the tactician with up to date advice on buoy type / RF channel availability (current and against plan), launcher load and any load changes required to satisfy tactical requirements. The system would also indicate the order of buoy drop.

## The Management of Sonics Information

The management of sonics information would be simplified if graphical or colour indicators were given to selectively assist in at least the following:

- Indication of information incorporated in the TT;
- Indication of information not associated with the target;
- Indication of accuracy of information when requested;
- Indication of 'stale' information.

Much of the above could be handled by an KBS using rules: 1) assigned and subsequently tuned by the tactician in the 'light of reality' and assessment of TT performance; 2) associated with the management of information utilised by the tracker.

## The Management of TT Utilised Information

The main problems associated with the management of TT utilised sonics data, and derived information, are associated with the adding of new sonics followers (sonics system related aids to keep track of designated sonics signals and associated data / information) into the TT and the deletion of old or rogue followers. The exercise involves the examination of the inevitable lists and the selection of list information - a time consuming exercise for the tactician. Often the indication that a list needs attention is given by graphical indicators of sonobuoy fixing or TT performance.

Simple rules could be devised to allow the automatic incorporation and removal of followers from a TT utilising KBS techniques. The rules would consider such as current TT position and performance, assigned accuracy of followers and correlation of the follower performance with that of the TT.

It is envisaged that the tactician would still need to tune the KBS. However, a work required to tune a system would often save the tactician unwanted workload at times of stress e.g. the few minutes before an attack on the target is performed.

## The Assessment of TT Performance

The tactician will be continually assessing TT performance in practice. The problem with current assessments is that they are time consuming and demanding on the operator. The performance indicators of the TT are frequently difficult to interpret and even more difficult to equate to the reality of the tracking accuracy of the TT.

*The above problems feed back to affect decisions on how many buoys to use, what tactics to adopt, what sonobuoys to load into launchers and what can be managed in the time available (identify, delete or add) to improve TT performance* - (the positioning and orientation of the sonobuoy pattern is another but related topic).

Moreover, the perceived difference between the TT tracking accuracy and reality varies with the relative positions of the target and sonobuoys and any changes in target performance. This is mainly due to definable errors that can be related to either buoy positioning, buoy type or target sound radiation patterns. Thus, KBS could be sensibly applied to alleviate the problem of TT performance assessment.

Furthermore, improvements to standard graphical depictions of the TT, to assist the tactician in accuracy assessment and indications of methods of amelioration, are not difficult. Figure Two illustrates one method of graphical assistance. There follows a brief discussion on the form of two possible KBS solutions to the assist in the continual determination of best target position.

## Tracker Selector KBS

This system would compare and select the optimum TT from several trackers working under differing interpretations of sonobuoy data but using appropriate rules concerning their efficacy considering target performance and sonobuoy relative positioning to the target.

The system would elicit information from the tactician when required (e.g. initialisation or on the approach to rule boundaries) and give plain language and graphical depictions of the system's conclusions, and where applicable and requested, the reasoning used.

## Tracker Tuning KBS

This KBS approach would allow the internal assessment and tuning of the rules employed in an associated set of trackers forming a TT. A KBS solution would be used to compare inputs of 'real world' target object (from operator inputs or through simulation on the ground) against each trackers distorted view of the world. The system

would heuristically explore the set of the trackers' rules in search for optimal performance.

The method could be used to: 1) progressively refine the set of algorithm rules used in the trackers; 2) Produce an intimate knowledge of the strengths and weaknesses of the trackers operating under optimal rule sets.
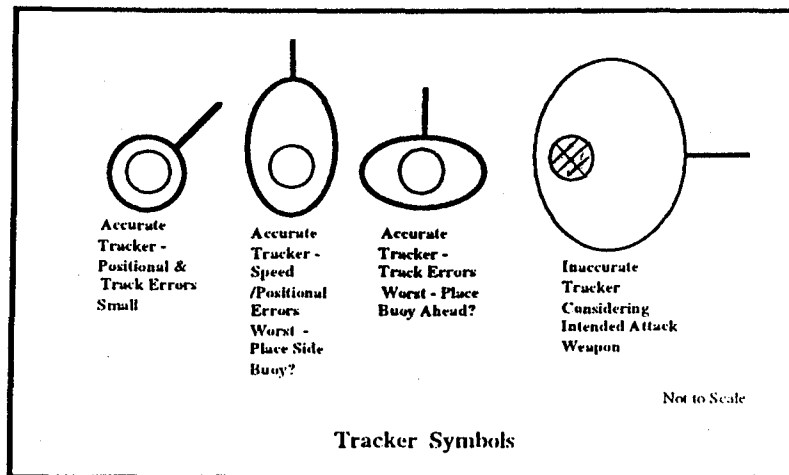


Figure Two - Possible set of TT Symbols Showing Fixing Error Forms and Accuracy of Information
(Darkly shaded circle to right given to indicate that TT information is too inaccurate to support an attack with the selected weapon)

## What Should be Achievable through Effective Adoption of KBS Solutions to this Problem Area

The following should be achievable:

i. The operator must be led to believe that known TT performance limitations are minor considering the uncertainties of the environment.

ii. The management of sonobuoy usage and associated TT data must be assisted and incur little unwanted workload.

iii. The TT should be capable of automatically adopting its assessed best tracking mode. The tactician should have an option to manually select a tracking mode.

iv. The presented results of the tracker must allow easy and unambiguous appreciation by the tactician.

v. Tactician's queries on tracker reasoning should be met by answers requiring little interpretation and commensurate with expected operator skill.

vi. The tactician should be allowed to tune the TT, possibly through a series of real world 'fixes. These fixes should be accepted by the KBS considering the TT rules evoked at the period of the fix.

## Final Comment

Most of the suggested areas that could be improved with the aid of KBS are areas where the work is easily definable but very time consuming and onerous to a human. It is a wonder that some form of 'rule based' assistance has not been

applied before in this area, excluding the use of home made china graph boards. The amelioration of workload is essential to allow the maritime tactician as much time as possible to cogitate on his working environment, the possible target intentions and the optimum tactics to be adopted.

The suggestions for improving the trackers are slightly more complicated, but not much more. The main adverse effects on tracker performance can be easily determined by considering the form of information available, the positioning and type of sources of information (sonobuoys) with relation to the tracked target and the target performance.

The uncertainties produced through the nature of the environment will always be with us. The aim with TTs should be to strive for a performance that is nearly up to the tracking performance possible from a skilled human. Only then will the skilled human believe that the assistance is worthwhile.

## References:

1. Adelman, L. (1992) *Evaluating Decision Support and Expert Systems,* John Wiley & Sons, New York.

2. Joint Acoustic and Anti-Submarine Warfare Centre, Sonobuoy Reference Card (latest edition).

3. Rasmussen, J (1986) *Information processing and human-machine interaction: An approach to cognitive engineering.* Amsterdam, North Holland.

# AIDING WEAPON DELIVERY

Douglas B. Hall
Delex Systems Inc.
5100 Springfield Rd. Suite 200
Dayton OH, 45431, USA

## Summary

In the restrictive rules of engagement inherent in the post cold war climate, tactical pilots must meet stringent collateral damage constraints during the delivery of air-to-surface weapons. This is exacerbated by the effect of employing sophisticated standoff weapons, many of which incorporate complex lock on after launch avionics. Determination of weapon delivery outcomes as a function of pre-planned mission data, captive carry air crew interactions and freeflight Man-in-the-loop control is an area where the electronic crewman plays a significant role. Further, while standoff weapons enhance aircraft survivability, it has restricted the use of traditional own aircraft sensors in pilot weapon release decisions. Programs such as TALON SWORD and TALON LANCE are demonstrating the ability to bring third party and national asset targeting information into the cockpit for pilot use, at the expense of significantly increased processing and information assimilation workload. How can we best incorporate aircraft avionics to digest the volume of available information, update standoff weapon mission planning data, and determine target acquisition probability for these complex weapon systems? What amount of further sophistication should be pursued as we define the next generation of aircraft, avionics and weapons?

## 1. Introduction: Aspects of the Current Situation which Affect the Role of Aircraft Avionics with Regard to Precision Guided Munitions (PGMs) Delivery

Military pilots today in performing the war fighting mission are faced with several factors that combine to create a significant rise in pilot workload. These include:

- An operational environment with restrictive Rules of Engagement (ROE) that prohibit collateral damage in the target area and exclude engagement of unintended targets. Damage to non-military targets and casualties to civilian personnel are often career ending errors. These ROE exist largely because of public and political perception that "high tech" weapon systems have the capability to prevent undesirable outcomes if properly used. Observe both the media treatment of weapon bulls-eye imagery from Desert Storm and the punishment of "operator errors" such as the shootdown of the Iranian Airliner by AEGIS and the recent F-15E shootdown of the Blackhawk helicopters in Iraq. There is an expectation that the men and machines are capable of surgical strike in virtually all circumstances, if the weapon systems are properly employed. This undoubtedly contributes to pilot stress during peak workload events such as weapon targeting and release sequences.

- Recent events in the international arms sales market. This results in a substantial increase in the "Gray Threat" as many more countries purchase western or ex-Soviet technology air-to-air (AAM) and surface-to-air (SAM) weapon systems. There will be in an increasing level of lethal air defenses that the pilot will have to contend with. If Desert Storm had been conducted in 1998, it is perfectly possible that the air defenses could have included SA-10 or later SAMs. The effect of an increased threat environment directly adds to pilot workload in the management of countermeasure systems, out of cockpit scan, and avoidance maneuvering requirements.

- The increasing sophistication of the current generation and projected PGMs. This increased sophistication has added accuracy, standoff and flexibility at the cost of increased Command and Launch System (CLS) complexity and large data processing / data management requirements. Intensive mission planning requirements now exist for operational systems such as the Standoff Land Attack Missile (SLAM)[1], and developmental systems such as Conventional Standoff Attack Missile (CASOM)[2], Arme de Precision Tres Grande Portee (APTGP)[3], Joint Standoff Weapon (JSOW)[4] and others. Aircrew are being required to understand and effectively use weapons that do not lend themselves to simple or intuitive launch processes. Further, while standoff weapons enhance aircraft survivability as the launch platform shoots

27

farther from the point defense systems of the target area, it has restricted the use of traditional own aircraft sensors in pilot weapon release decisions. Many of these weapons require the pilot to serve as a primary or back-up source of target identification through imagery transmitted by the weapon. This task is mission critical, time limited, and workload intensive.

• A significant shift in the operational concept to integrate off-board sensor system data from ELINT, imaging satellites, targeting and control platforms such as Joint Surveillance Target Attack Radar Systems (Joint STARS), with onboard aircraft sensors. Programs under the Tactical Exploitation of National Capabilities (TENCAP) such as TALON SHOOTER and TALON VISION are engaged in the research and development of systems and processes that place this here-to-fore highly restricted information directly in the hands of the warfighter.[5] This near real time satellite and third party targeting data, when integrated into the cockpit, will provide to the pilot an order of magnitude leap in the quality of data needed to increase situational awareness, with the expense of unprecedented increases in the the data processing and information assimilation workload.

How can we best incorporate aircraft avionics to digest the volume of available information and update PGM mission planning data? How do we structure the role of the electronic crewman to reduce pilot stress, tasking, and workload, while improving the predictability of weapon delivery and reducing undesirable outcomes such as collateral damage? The military aircraft industry stands at the brink of a revolution in the capability of aircraft avionics to undertake new roles, especially regarding the potential for an order of magnitude increase aircraft mission computer and stores management system processor power and memory capacity. It is sobering to consider that current generation tactical aircraft have, on average, processor power of approximately 1 Million instructions per second (MIPS) with a nominal 1 megabyte of memory, while off the shelf processors of 300 MIPS and 1 Gigabyte memory capacity can be obtained in a portable personal computer. The remainder of this discussion will examine three areas where the electronic crewman is under utilized, but where it can have a significant effect on reducing pilot workload and increasing weapon delivery effectiveness.

## 2. PGM Mission Data Manipulation in a Changing Threat Environment

Aircrews today have limited access to real time changes in the threat environment that formed the basis for mission planning. With current systems, even if the aircrew is provided with timely, accurate real time threat data, he is left to intuitive processes to effectively react to the new information. Current generation aircraft do not automatically integrate new threat data into the aircraft or weapon mission plan resident in the aircraft avionics. Further, current generation PGMs do not lend themselves to optimized delivery through intuitive processes. Many PGMs require precise, well-defined launch envelopes and extensive, data-intensive mission plans that cannot be derived intuitively. With current CLS designs, reacting to real time threat updates requires tedious and time consuming manual keypad inputs of new launch point and weapon route of flight data within the mission plan. Very often this type of manual PGM mission re-planning is not possible in a fluid, high threat environment where time-on-target (TOT) and other mission constraints may limit the time available for re-planning to only a few minutes or less.

Even with PGMs, such as HARM, that do not require extensive mission planning inputs, the aircrew must still determine how to best position the aircraft for launch in a high threat environment.[6] Currently, when he receives a new threat target, the Suppression of Enemy Air Defenses (SEAD) aircrew must intuitively determine the most survivable area from which to launch with respect to the threat, and how to best navigate his aircraft to that launch envelope.

The often strict limitations on employment of PGMs make it very difficult for the aircrew to effectively react to new information. Because he cannot manually re-plan his weapons' delivery, or simply because he cannot intuitively determine a survivable route to a new launch point, the aircrew currently has the poor choices of sacrificing survivability in an unplanned threat environment, delivering his ordnance in a degraded mode, or aborting the mission altogether. These tactical employment constraints result from limitations in aircraft tactical communications, computer memory, and processor speed. Advances in these areas, however, now make it possible to store and

process real time threat data, and to apply that data to update aircraft and weapon mission plans for increased survivability and effectiveness. An approach to fulfilling the needed capability is the integration of an onboard mission planning or replanning module. Describe below is a concept for operation (see Figure 1) for such a system.

• Mission Planning. Aircrew plan the tactical mission using systems such as the U.S. Navy Tactical Mission Planning System (TAMPS) or Air Force Mission Support System (AFMSS). During the planning phase, a database including current threat beddown data, Digital Terrain Elevation Data (DTED), weapon characteristics such as turn and climb rates, aircraft characteristics and weapon delivery parameters is created. This database will then be inserted into the aircraft's data transfer device (DTD) along with the aircraft and weapon mission plans. Table 1 is a listing of the projected data generated during mission planning.

**Table 1  Mission Planning Data**

| Data Description | Characteristics |
|---|---|
| Aircraft Data | Route Lat/Lon/Alt, fuel projection, event description, mission timing, etc. |
| Weapon Data | Route Lat/Lon/Alt, fuel projection, event description, mission timing, etc. |
| Area Data | DTED, Probability Vertical Obstruction Data (PVOD), no fly zones, etc. |
| Threat Data | Location, Threat Identification Number (TIN), threat status, uncertainty, CM (time), etc. |
| Signature Data | RCS, threat effects |
| Aircraft ECM | ECM and ESM cueing data, ECM effectiveness data |



Figure 1  Inflight Route Replanning Operational Concept

29

- Preflight Phase. In the preflight phase the aircrew will load the mission data for the aircraft and weapon in a single data transfer operation using the aircraft DTD.

- Employment Phase. In the employment phase, the aircraft enroute to the target receives offboard threat updates through an aircraft data link such as the Joint Tactical Information Distribution System (JTIDS), Tactical Data Information Link - J (TADIL-J)[7] or some other dedicated link capable of providing real-time updates to the threat beddown. The aircraft avionics merges offboard and onboard (own sensor) data[8] and uses the resulting data in a route evaluation process. In this process the new threat beddown information is used to assess the resulting impact on the aircraft route of flight, weapon launch envelope and weapon route of flight. If the outcome of this assessment is a determination that survivability along the pre-planned route of flight is jeopardized beyond a reasonable threshold, it will cue the aircrew that the mission route may no longer be viable.

The aircrew then may request a display of the resulting aircraft route of flight options, PGM launch envelope options, and weapon route of flight options appropriate to the updated threat beddown. The options are based upon threat data, aircraft/weapon performance data, signature characteristics such as radar cross section (RCS), airspeed, and ECM suite data. The generated route and launch zone options are optimized for aircraft and weapon survivability, (within constraints set during the mission planning phase such as minimum altitude, no-fly zones, time on target, and fuel on-board), offering the aircrew several aircraft route, launch point, and weapon route options as appropriate to the updated threat situation.

When the aircrew selects an option, aircraft route and launch point data is formatted and output as a change to the mission plan in the aircraft mission computer. The aircrew will then receive new steering and navigation cues through the normal flight director functions in the aircraft OFP. If a new PGM weapon route of flight or launch point is required, then the electronic crewman formats the data as required for the weapon and passes the updated weapon mission plan to the weapon through the weapon data bus such as the MIL-STD-1553 bus interface. This system can also offer the aircrew other selectable options such as a survivable route of flight from the threat zone to either home base, a tanker track, or other designated point. This capability will greatly improve current and future weapon system effectiveness by providing increased aircraft and weapon survivability, and weapon delivery optimization. Mission flexibility also increases with the capability to redirect airborne aircraft to new high priority targets. For example, an aircraft on a Battlefield Air Interdiction (BAI) mission employing JSOW could be redirected to a high priority SEAD mission by passing the new threat coordinates to the aircraft and allowing the electronic crewmember to compute a new mission plan. Currently this task is done in an error prone manner requires an inordinate amount of time and computation on the part of the aircrew that is not generally feasible in fighter type aircraft. Aircraft and weapon survivability will be enhanced by redirecting their routes of flight around new threats not included in the original mission plan. The new routes of flight will maximize survivability allowing the aircrew to accomplish the mission with minimum threat exposure. This capability will be particularly effective in the SEAD role where target locations are often not known in advance.

## 3. Use of Aircraft Avionics to Predict Weapon Outcomes

Lock on after launch weapons such as HARM and HARPOON[9] select targets as a function of complex processes that are not intuitively obvious to the pilot. To further complicate the problem they have multiple modes of operation that vary the behavior of the weapon in the target selection process. For example, with HARPOON, the pilot designates a target location or bearing and the missile proceeds along a flight path, initiates a seeker search to acquire the target, and following detection, tracks the target to impact. Several options, including missile search mode, missile function during search, the use of waypoints and the location of the missile targeting solution (target position), combine to determine which target the missile finds and ultimately impacts. For example there are tactical situations where multiple ships or land is in the target area, the pilot selection of missile mode and targeting solution are critical to the outcome of a HARPOON engagement.

The above example also highlights another aspect of the problem; the use of complex weapons in scenarios that are on the

boundary of the design specifications. HARPOON, for example is designed to attack a single ship target in an open ocean environment, however, the missile has excellent capability in more complex scenarios, providing the pilot chooses the optimum targeting solution.

One of the solutions used by command decision makers in setting ROE for use of weapons like HARPOON to ensure outcomes is to restrict standoff. It is noteworthy that the use of HARPOON in the Gulf of Sidra against Libyan Navy vessels required visual identification by over flight and launch ranges were minimized.

For weapons such as HARM the seeker footprint is a function of mode and launch range. This is central to target selection, and is currently left to the ability of the pilot to intuitively determine the "heart of the envelope". The ability of the pilot to effectively perform this task is related to how well he can recall complex charts from the Tactical Manual.

A solution for this problem is to use an aircraft processor to evaluate the targeting alternatives and to present the pilot with optimum solutions. The type of weapon system modeling necessary is similar in nature to engineering and training simulations Delex Systems has produced. This type of simulation requires detailed algorithms for weapon sensor performance, weapon flyout dynamics, and will most especially require a method of entering the current target environment to define the scenario for the simulation. The type of mission database and the activity discussed in section 2 above could perform this function. The output of the simulation is some numerical non-subjective determination of the weapon probability of acquisition ($P_{acq}$) of the intended target and an assessment of the probability of engaging a non-combatant or unintended target.

The effect of threat system performance on weapon probability of arrival would increase the fidelity of this process and could generate actual probability of arrival projections.

## 4. Use of Weapon Imagery in the Cockpit for Man-in-the-Loop (MITL) Target Designation

Weapon Systems such as WALLEYE, AGM-130, SLAM, CASOM and JSOW (P$^3$I) use a data link system which supports the weapon acquisition of the target during the terminal phase of weapon flight. This process, critical to the mission success, is one of the highest workload items a pilot undertakes. The best example of the difficulty of this task can be found from use of the SLAM in Desert Storm during a SLAM mission a pilot tried to complete the target designation task while simultaneously maneuvering to avoid an incoming SAM and handle an illuminated master caution light[1].

Two aspects of the target designation task, where the electronic crewman is central to aiding in this mission critical process, are target recognition and hitpoint designation.

The first phase of the task is target recognition. In this task the pilot views the target area image sent from the weapon and determines the location the target in the image. Several aspects of the current design of such systems can be improved with the application of current technology.

Transition from the analog video link systems to the use of digital video will improve image quality and allow application of image enhancement algorithms. This type of image enhancement will be especially useful for weapon systems applied under boundary conditions where sensor images are degraded due to weather, countermeasures, or image source limitations such as minimal target thermal variance for imaging infrared sensors.

An obvious improvement would be the integration of a larger, high resolution display in the cockpit. However, space limitations makes this simple solution difficult to achieve. An alternative is to integrate a system that could allow image magnification.

Target recognition can also be improved through the integration of target imagery into mission data so that the pilot can have a reference image for comparison with the weapon imagery. Such imagery in the cockpit is currently in hard copy form, which has obvious handling limitations.

The ability to pass imagery and location data to the pilot for application of these imaging standoff weapons against mobile high value targets such as SCUD launchers is a fundamental aspiration of the TENCAP projects[10]. The role of the electronic crewman in automated processing and effective display of this data is essential. In the integration process of these

31

capabilities the importance of target cueing must be central.

The second phase of the target designation task involves the mechanics of executing the hitpoint designation procedure. In current data link design and aircraft integration, the weapon data link system are generally carried externally in conjunction with the weapon stores on a wing station. This integration mechanism results in a procedure that is generally as follows:

• The pilot completes the target recognition process and initiates hitpoint designation with some cockpit switch

• The initiation command is sent to the weapon where a cursor is inserted into the target imagery

• The pilot sees the cursor and initiates some cursor slew mechanism

• The aircraft avionics reads the slew commands, formats it for transmission to the data link pod

• The data kink pod receives the slew information, changes the data into the command link format and transmits it to the weapon

• The weapon data link receives the message and decodes it, and passes the command to the targeting sensor to move the cursor on the internal weapon data bus

• The targeting sensor moves the cursor in the image

• The image is passed to the weapon data link for transmission along an internal weapon video bus

• The weapon data link transmits the image to the aircraft

• The aircraft data link pod receives the video and sends it to the cockpit for display

• The pilot sees the cursor move in the image and adjusts to bring the cursor over the desired hitpoint

• The pilot hits a cockpit switch to "designate" the hitpoint.

This type of multi-step process can cause significant time delays between pilot action and movement of the cursor in the display. A more direct integration of these systems with the aircraft avionics, where the hitpoint designation process involves fewer steps, data hand-offs and reformatting will greatly reduce these delays.

In addition, in conjunction with the previous recommendation for integration of a larger display to support target recognition, the inclusion of a touch screen or other one step hitpoint designation is possible.

## 5. Conclusion

There are a number of areas where the electronic crewman can significantly improve the effectiveness of delivery of PGMs, and simultaneously reduce the pilot's workload. Especially important will be the approach that industry takes with respect to enhancements in aircraft capability regarding processor throughput and memory capacity, integration of offboard threat information and target imagery, and improvement in target recognition and designation processes for MITL weapons.

## 6. References

1. Kandebo S.W., "Upgraded SLAM", AW&ST, 21 March 94, pp 53-54

2. Shifrin C.A & Morrocco J.D., "UK, US Firms Vie for CASOM Contract", AW&ST, 12 Sept. 94, pp 73

3. Unk., "French Enter Land Attack Conventional Cruise Missile Arena", ASD, ISSN 0193-4546, 27 May 94, pp 328

4. Unk., "Desire to Accelerate JSOW Rules Out some Seeker Technologies", ASD, 23 May 94, pp 291-292

5. Scott W.B., "Space Warfare Center Supports Warfighters", AW&ST, 28 Mar 94, pp 64-65

6. Unk, "Jane's Weapons Systems", 1988-89, 1769.311, pp 724-725

7. Unk, "Concept Papers for Planning Purposes on $C^3I$ for Multi-Role Fighter", CBD PSA-0743, 16 December 94

8. Unk, "Expanded Situational Awareness Insertion", CBD, 1 March 94

9. Unk, "Harpoon and SLAM Weapon Systems", MDMSC, January 91

10. Munro N., "Demonstrations Propel USAF Weapons Efforts", DN, 5-11 July 94, pp 8

# CAMA: Some Aspects of a Military Crew Assistant System

Brugger E. and Hertweck H.

DASA-LM, LME-Friedrichshafen/Munich, Germany

## General

CAMA, an acronym for Crew Assistant Military Aircraft, is a joint research and development program of the German Aerospace Industry and research establishments.

The program CAMA, in principle foreseen for all type of aircraft, has currently a priority in a two pilot cockpit as it exists in a transport aircraft (fig.1). The system will assist the crew in an enhancement of the situation awareness in all mission phases from starting flight planning by the crew until to the debriefing after landing. It provides defined types of operational support, e. g. preparation and execution of tactical navigation or critical loading effects on a restart from an unprepared field. Basic services are available as check list routine work, fuel/load management and so on. One of the main task however is a comprehensive capability for alert and warnings in different way at the crew-aircraft-interface using ability of seeing, hearing and tactile sense.

The worksharing in the program, which is currently under contract of the German BWB (Agency of Procurement) consists on the industrial side of the military aircraft division (LM) of the Deutsche Aerospace and the Electronik System Gesellschaft (ESG) München and on the research side of the Institute of System Dynamics and Flight Mechanics of the Universität der Bundeswehr München and the Flight Test Center of DLR (German Aerospace Establishment), at Braunschweig.

This cooperation is seen as an excellent opportunity to combine the latest results of research and development on the field of cockpit assistant systems and the experience of military operational requirements in future air weapon systems.

The question "why CAMA in current and future military crew stations" reflects the today situation in the cockpit of modern weapon systems. A lot of R & D has been performed in this area. Therefore, only some spotlights will be given with respect to this complex.

Facts are:
- All pilot actions are canalized by a given MMI, which is a result of the chosen cockpit-layout
- All decisions are made by the CREW as a result of mental (more or less routine) processes
- Decisions effected faulty by the CREW are recapitulated by the system
- During the decision process in the cockpit the critical point is the CREW
- In a highly automated cockpit, the pilot can find again himself in a situation Out-of-the-Loop

An approach to correct this situation is represented by CAMA.

What is CAMA doing?

The system

- Acquires continuously aircraft data, mission data, environment data and pilot's data
- has access to stored data: static data base, dynamic data base and knowledge based data
- generates situation dependend on recommendations for pilot actions
- makes available knowledge based, static and dynamic data to the crew
- executes planning tasks
- does not make autonomous decisions and does not intervene to the weapon system **without** crew order
- does not increase the automation level of the weapon system

The program CAMA consists of a number of 16 separate modules for data acquisition and a further module for data control (fig.2).

These are:

**Low Altitude Planner**
It provides a low altitude flight path to the target using the terrain profile in order to minimize threat; optimize flight parameters and fuel management

**Terrain Interpreter**
Comparison of flight performance envelope with obstacles in low level flight, recommends possible evading manoeuvres.

**Computer vision outside**
Interpretation of optical sequences to recognize the relative present position and obstacles.

**Computer Vision Inside**
Acquires sequences and interpretes the crew movement with respect to body, head and eye movement of the for identification of crew actions

**System Interpreter**
Determination of the status of the a/c-systems: data supply for other modules

**Environmental Interpreter**
Acquisition, monitoring and recording of meteorological data

**Tactical Situation Interpreter**
Acquisitio,presentation and recording of the tactical situation, threat during the whole mission

**Flight Situation - and Threat Interpreter**
Acquisition, recording and assessment of data, which will be interpreted as "conflict situation". Specifies a conflict situation

**Pilot Behaviour Reference**
Creating a crew model related on a "normative behaviour" of the crew during the whole mission
Knowledge base with respect to the margins of allowable pilot behaviour

**Pilot Behaviour Deviation**
Creating a crew model for deviation of the reference behaviour

**Pilot Intention/Error recognition**
Continuous comparison of reference behaviour with factual crew behaviour.If there is a deviating from the mission plan, it will be examined, wether there is a true intention or a pilot error existing.

**Flight Status Interpreter**
Examination if and at which point an expected flight progress has been achieved

**Automatic Flight Planner**
Continuous updating of the mission plan on the basis of current change of tactical and operational situation

**Aircraft Interface**
The interface picks up all necessary data of the a/c systems (and feeds the a/c-systems

34

with CAMA data) (this part is under investigation, the data transfer from CAMA to the a/c has to be subject to a hazard assessment).

**Dialogue Manager**
It is the CAMA-CREW Interface established by display, direct voice input, direct voice output and manual control.

**External Communication Interface**
It is the input for all arriving, transceiving data via data link for mission relevant processing.

The structure of CAMA interface foresees a core part for all CAMA-relevant computations. The communication with crew, aircraft, outside world occurs by interfaces for CAMA-specific sensor data, aircraft specific informations, data link informations and dialogue with the crew.

They are embedded in the overall CAMA-structure. Fig. 3 shows this struture with respect to the tools used during this phase. It gives a functional overview on all CAMA modules and data sources including the data flow among them. The modules may be grouped into the following seven classes

- the CAMA relevant aircraft systems
- the situational data generation modules
- a data pool describing the actual situation
- a data pool for reference data
- the modules for analysing, planning and conflict solving
- the CAMA MMI
- modules for pictoral data processing

The basic functions of each module have been described above, however, a few words remain to be said about the data bases at the bottom of the

figure. They contain the knowledge CAMA advices are based on. There are on one hand data describing physically existing outside world items such as navigation aids, terrain elevation and feature data, on the other hand modelling data for the simulation of a reference pilot and a reference aircraft.

After this overview, for some of the modules under investigation at DASA their functions will be explained in more detail.

**Terrain Interpreter (TI)**
The terrain interpreter monitors the CAMA aircraft's flight path and informs the pilot and the CAMA system whenever a risk of flying into terrain occurs. This is done by continuously calculating a three dimensional performance envelope of the aircraft (fig. 4) taking into account the current state of flight and the actual aircraft configuration. The performance envelope spans between the flight path at minimum turn radius for level flight and the flight path flyable with zero turn rate and maximum climb performance. Whenever this envelope collides with the terrain as defined by the digital data base the respective flight pathes are prohibited and depending on the urgency of action an immediate pilot activity is commanded or a replaning process is triggered.

**External Communication Interface (ECI)**
The module external communication interface simulates digital data-links between the CAMA-A/C and any possible outside world communication systems.
The operator at the ECI-station initiates and change the scenario in

the operating area of a CAMA-mission. For a efficient work the ECI implements a graphical user interface based on the X-Windows- standard and OSF-motiv. Fig. 5 shows an example.

Functions of the ECI are:
- setting or changing elements of the tactical scenario such as threat or wepon engagement zone.
- setting the status of an airport or navigation equipment.
- setting or changing wether conditions in the operating area.
- changing or preparing a new mission plan for the CAMA-A/C.

**The Crew Assistant Graphical Interface Shell (CAGIS)**
This module allows manual crew inputs to the CAMA system via display and a designator (e.g. joystick, trackball). The CAGIS screen is devided into four sections: The lower part of the screen displays at all turns the basic menue, the right half portion above it offers submenues to the selected basic main functions, the upper line shows the status of the input and on the rest of the screen a map of variable size is displayed allowing e.g. selection of NAV stations, setting flight plan data, sending ATC commands, trigger service functions, operate aircraft systems. It is designed to avoid the requirement for any alphanumeric input device. A special purpose is seen in situations, where data are to be altered by the crew due to actual mission conditions which are not covered/updated by the CAMA missionplanner.

**Development Status**
Currently all modules are functionally specified and some of them are already coded and have underwent initial testing. Discussions between the partners have shown that refinement will be needed e.g. to avoid functional overlapping between the modules and already existing aircraft monitoring systems. Starting approximately in April '95 the modules will be partially integrated together and an overall CAMA simulation is envisaged beginning in 1996.

The simulation shall test a basic CAMA-version. After successfull checks simulator flights by different pilots are planned. The system is foreseen to enter flight tests with the ATTAS (Advanced Technologies Testing Aircraft System) of the DLR in about two or three years.

Situation
Awareness

Alert
Warnings

Basic
Services

Operational
Support

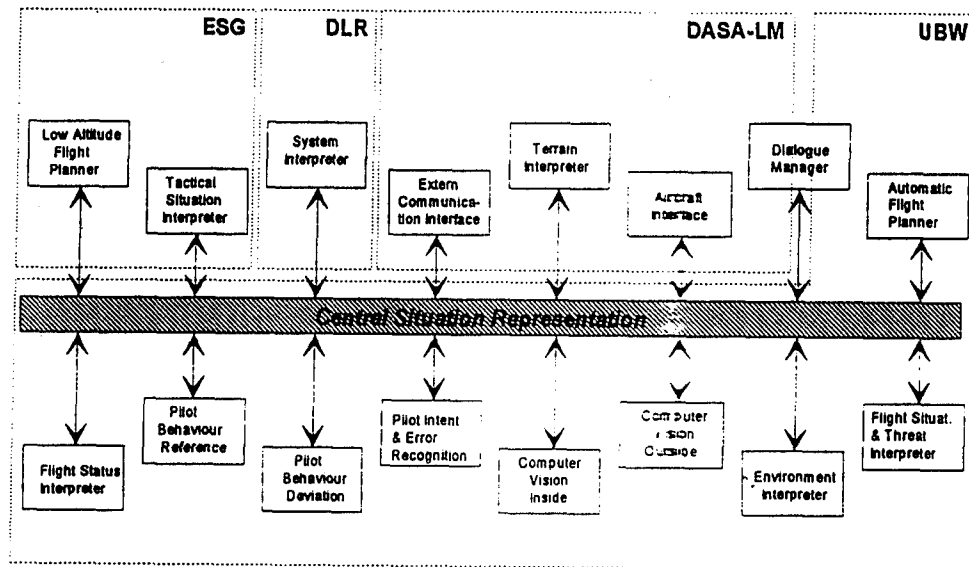Fig. 1: The System CAMA



Fig. 2: CAMA Modules
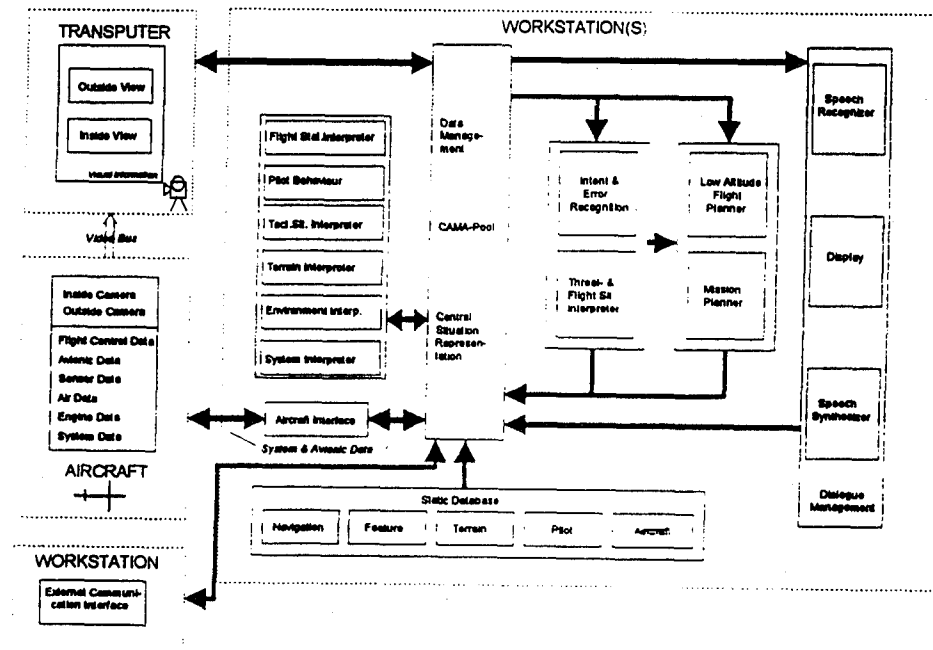


Fig. 3: CAMA Structure

Fig. 4 Flight Performance Area



Fig. 5 Graphical User Interface of the modul ECI

# BATTLE-SUITABLE, ELECTRONICALLY PROVIDED INFORMATION

J.S. Seeman
McDonnell Douglas Helicopter Systems
5000 East McDowell Road
Mesa, Arizona, 85215, USA

&

T.R. Metzler
U.S. Army Aviation and Troop Command
St. Louis, Missouri, USA

## SUMMARY

This paper provides a user-based orientation useful to assist in understanding, organizing and designing future, responsive electronic-based decision-making aids. It argues that electronic decision-assisting systems, using fuzzy logic to overcome human limitations resulting from selective perception and risk minimization, may be more valuable to management of battlefield information (could provide more, quality data, or information) than currently thought. It proposes that such systems be considered not only to support cognitive decision-aiding in future aircraft crew stations but also perhaps to autonomously make and implement the results of such decisions for all players, at all levels of the battlefield.

## 1. INTRODUCTION

Current trends for the future battlefield appear to be leading toward proliferation of electronically provided information. The purpose of such information is to improve the effectiveness of combat operations. Much costly effort has, is being and will be expended to accomplish that purpose. Properly implemented, this technological advance can prove as powerful a factor to success in future battle operations as the introduction of gunpowder.

This paper is an attempt at introduction of considerations which may bear upon the ultimate effectiveness of the electronic information management systems currently in development. The fundamental approach taken here is that electronically provided information is likely to be useful to combatants to the extent that it contains advice or guidance relevant to their immediate needs.

Several characteristics are believed to be important to the nature of the information intended to improve the effectiveness of such operations These include speed of response to changing events, recognition of options available for action, prediction of the implications (in terms of future capabilities of friend and enemy) of adoption of a particular option or option set, and knowledge of the consequences of utilization of the available options.

Relevance of information to immediate needs is a function of its accuracy and timeliness as well as its compatibility with the user's cognitive framework. Taking such considerations into account will result in providing information suitable to the user. Neglect of these factors will merely add to their response time lag. Selected approaches to improving information relevance for the various levels of information management within the battlefield, are discussed.

## 2. CONCEPT DESCRIPTION

### 2.1 Assumptions

For purposes of this paper, factors relating to the collection and dissemination of information through electronic means are assumed to have been satisfactorily resolved. Thus, appropriate sensors, communications channels, communications covertness, range satisfaction, data fusion and validation, and verification of accuracy of data transmission are all considered to be in place and operating satisfactorily.

Most data collected may appear to be certain, or non-probabilistic, but in most cases will be uncertain. Thus, data should be treated as probabilistic. If probabilistic data is forced into "binary certainty" format, it will become a constant source of unquantifiable error.

It is also assumed that all the information on the battlefield is provided for application in as timely a fashion as possible. This is perhaps the one fundamentally most important factor to effective information use and known as significant challenge to the designers of the digitized battlefield.

Further assumed is that available data will potentially be used by human operators at all levels involved in prosecution of the battle.

Available for their use will be the plethora of modern information including: Global Positioning System Data, Digital Maps, Weather Data, Intelligence from remote space, airborne and surface sources as well as data from less remote sources including proximal air and surface combatants as well as onboard sensor systems.

## 2.2 Background

For purposes of this paper, information users have been arbitrarily grouped into three major categories roughly corresponding with the organization expected during future battles. It is expected that the organization will include participants at essentially three levels: Theater Command Control and Communications (analogous to the Corporate Executive Level), users at the level of implementation (analogous to the Management Level and users at the action level (analogous to the Employee level). Unique information needs pertain to each level in the net.

The Executive Level requires data on trends of enemy and friendly activities to the extent that these relate to the accomplishment of established goals.

The Management Level requires data on details of accomplishment of subsets (defined by time and/or location) of the mission objectives and the likelihood of trend continuance.

The Employee Level requires data on the immediate objectives being pursued.

Various methods of improving the value of information to each of these levels of operation have been proposed. Those approaches can be grouped as the "game theory" methods. They essentially attempt to predict future events and conceptually test the efficacy of alternatives. Then, from among the alternatives, they attempt to recommend one or more which appear to have value in winning the game.

For use during battle, these approaches have not gained significant acceptance, probably because of human mistrust in their ability to appropriately take all relevant battlefield factors into account and also probably because of the belief that the safety of human lives in battle is best assured by human decision makers on the spot.

History and psychology suggest that both of these positions is suspect.

Historical evidence suggests that many battles have been lost because of information mismanagement by humans. Some battles almost certainly have been lost because some relevant available information was not taken into account. Conversely, some have been won because of

human ingenuity and the willingness to accept greater risk. Unfortunately, it is a certainty that human decisions during battle don't always result in preservation of the lives and resources of friendly combatants.

This state of affairs may be improved if humans recognize some of the major limitations which have repeatedly been pointed out by the field of psychology and try to overcome them through use of available technology.

These limitations reflect human shortcomings in the capacity to accept all available information from the world and manage it with dispassionate certainty. With the increased quantities and uncertainties of data which will accrue on the future digital battlefield, it will become more difficult for humans to formulate and extract a unified concept of the existing and emerging state of the battle. Otherwise stated, it will become more likely that humans will be the source of more erroneous decisions on the future battlefield.

Those human limitations which are expected to result in significant errors in judgment reflect the human characteristics of selective perception and predictive conservatism.

Selective perception is the well known phenomenon (perhaps one of the truisms in psychology) which is thought to result from individual human experience and which results in attending to data congruent with preconceptions or strong habit patterns, while ignoring those data which are incongruent. During battle, the ability to attend to all data is further limited by the exigencies of the battle. An excellent anecdotal example of this proclivity is found in the Japanese story known as Rashamon in which a revealing account is provided of the accuracy of eyewitness reports of a murderous attack.

Time stress and the emotional consequences of drawing potentially incorrect conclusions from data can also affect selective perception. Time stress in human processing is thought to affect perception by reducing the amount of data perceived. Consequences of decision making also can result in tendencies to reduce the certainties of decisions made, or inferences drawn, from available data.

Edwards, in the early to middle '60s, performed a series of experiments (summarized in Ref 1) using mathematical manipulations, based upon Bayesian probability theory, to make predictions of future real-world events. He compared these event predictions and their assigned confidence levels with event predictions made by humans using the same information provided for Bayesian analysis. His results can be said to have established that purely mathematical,

probability-based, methods can predict future events with greater associated certainty than humans. Unfortunately, these research data have not enjoyed wide dissemination.

If Edwards' research was correct, and was accomplished without the extensive use of the powerful computational systems that we are now examining for installation on aircraft and mobile ground vehicles, the availability of the newer computational systems we are currently projecting for use in the battlefield might possibly benefit from the findings of thirty years ago with resulting improvements in the use of probabilistic data.

Today, the U.S. Army is examining the potential combat advantages which might accrue if powerful airborne computers could assist in intelligently managing data currently managed by crew members alone. This approach has been applied to fixed wing missions on programs such as the U.S. Air Force Pilot Associate (P.A.) program. Other programs such as the Day/Night Adverse Weather Pilotage Program (D/NAPS) and the more recent Rotor craft Pilots Associate (RPA) program are examining in detail various approaches to gaining combat advantages for rotary wing attack and reconnaissance aircraft. Soon to come will be similar efforts to examine how computers might similarly assist tank crews. Additional efforts are directed at upper levels of command and the management of data for the Digital Battlefield. Further efforts are likely to be directed in all these areas toward increased and extended Battlefield Operational Capabilities.

These are all directed toward improvements in "Combat Effectiveness". If successful, they can be expected to result in improving the speed and accuracy of coordination of diverse inter and intra-service intelligence and combat resources. If these resources can effectively be integrated, the longtime dream of many practitioners of warfare may be realized; that is, the ability to simultaneously project concerted, diverse forces against an enemy toward a decisive end. Historically this approach has been called "Blitzkrieg". In more recent times it has been referred to with the phrase "Massive First Strike". Normally, however, such a strategy is associated with aggressive postures. It may also be viewed as a special defensive posture, consistent with reduced willingness to accept battle losses in a crisis response situation as was seen in Operation Desert Storm.

Preventing the achievement of this capability are numerous factors, not least of which are insufficient and stale data, erroneous data, and insufficient or unreliable communications These are the problems being addressed by developments in the digital battlefield as well as pilot associate programs. To the degree that the

efforts are successful, these problems are likely to diminish in importance.

These approaches are typical of those successfully undertaken in the past. They attempt to increase the quantity of precise, accurate data available for use by battle participants. This is today seen in information management thrusts intended to accomplish "horizontal integration" of information such that a "shared view" of the battle (situational awareness) is available to all management and combatant levels. While the horizontal integration approach of the impending digital battlefield is likely to improve the quality of data available for drawing inferences, it does not improve the quality or the speed of the inference-drawing process which we view as crucial. Not now being addressed by any existing program is how best to overcome the human limitations mentioned earlier.

## 3. PROPOSAL

It is desirable that these human limitations be surmounted at all battle levels: "Executive", "Management" and "Employee". While the fundamental techniques employed to overcome the human limitations mentioned earlier at each level of the battlefield hierarchy may be similar, differing products suitable for each level may be desirable.

The goal of this proposal is to provide quality data, or information, to those who require it. Quality data, or information, is raw material which has been processed. Processing should accomplish at least two objectives. It should first take into account all verified data, and then attempt not merely representation of the current state of affairs but also rapid prediction of future states. A nicety would consist of the conversion of future-state knowledge into action terminology, or orders intended to make best use of resources and to communicate that terminology to those affected.

It is perhaps appropriate to distinguish at this point between Data and Information. Data is defined as the raw material from which information is generated. Thus, data is the location of an enemy position or the type of weapon thought to be at that location. Information is the result of data manipulation, or processing, it results in unification of multiple data sources and provides guidance or direction for action.

Thus, to a battle commander, enemy intentions (direction and rate of movement, size and type of force), likely intercept or observation points and friendly forces available to counter enemy activity are information. In an aircraft, SAM location and status are pieces of data. Aircraft

41

response options (safest flight path, minimum signature attitude and countermeasures to be employed) in light of the threat activity is information. On the battlefield today humans are used to convert data into information.

The usefulness of the distinction between data and information and its relationship to overcoming human limitations at all levels for the electronic battlefield will become clearer as we progress.

Simply stated, during battle the Executive level attempts, through human reasoning, to interpret data available to it to discern and act upon trends (opportunities and hazards) toward accomplishment of established goals. The Management level, using its data sources and human reasoning, attempts to track, anticipate and manage battle contingencies as they arise. The Employee level attempts to track, anticipate and respond to management directives while minimizing battle cost/benefit ratios.

At each level, accuracy of data and speed of proper action are primary considerations. Yet, even if data is fresh and accurate, reliance upon the human in a data interpretation/decision-making role can be expected to be rife with error- from not taking all data into account and from drawing conservative or low risk conclusions. Additionally, the human decision-making process can be slow.

If it is assumed that the emerging electronic or digital battlefield can be expected to provide more accurate and timely information to all levels, the goal should clearly be to reduce reliance upon human decision-making.

What alternative to that process exists? The cavalier answer, of course is computers. A better answer is "smart" use of these computers to improve upon the process.

Humans tend to arrive at probabilistic decisions using a process of probabilistic reasoning. Fuzzy Logic as described by Zadeh (Ref 2), in its broadest sense, attempts to convert reasoning with uncertainty into computer manageable data. Among its goals are development of improved methods of using human knowledge-based systems and use of data banks to answer human queries about real-world events.

Use of approximate reasoning approaches to management of data on the battlefield offer the possibility of more powerful use of the available data. These methods, compared with human beings, are capable of rapidly dealing with all the data available using a technique approximating human reasoning, but without predictive conservatism. Thus, predictions resulting from their use should more accurately represent future

events (be associated with higher willingness to accept risk) than human predictions.

Bhatnagar & Kanal, (Ref 3) suggest in their paper on methodologies useful in approximate reasoning that "...a battlefield planner may be interested in hypothesizing an enemy's plan of attack" which may be more useful than knowing the "probability of being attacked by the enemy."

In Zadeh and Kacprzyk, (Ref 4), detailed papers are presented on available and emerging techniques of logically dealing with the extraction of information from uncertain data using purely probabilistic as well as hybrid methodologies.

Such systems can be expected to solve the limitations of human processing and simultaneously extract more certain information from uncertain data. Fuzzy Logic can help in the conduct of both steps in the generation of information-unification and action guidance. However, fuzzy logic alone may be insufficient for resolution of all information management and decision-making problems. The primary limitation to fuzzy logic may be its reliance upon humans and the process of knowledge acquisition to devise the expert systems of which it will consist. For fully responsive systems to evolve, it will perhaps be necessary to explore the construction of hybrid systems comprising fuzzy logic, purely binary as well as Bayesian algorithms.

Some general and specific examples may help visualize the possibilities of such systems.

Recall that it has been argued that at all battlefield levels the human is principally relied upon to integrate data and predict future events from current data. Once the events have been predicted it is of course necessary to convert knowledge of the events into a course of battle action, formulate and allocate orders and transmit these orders to implementers.

At the "Executive" level, the improved digital battlefield data can be automatically formatted for processing by expert-fuzzy logic systems whose objective would be to identify and predict future battlefield trends, formulate plans to maximize friendly success probabilities, convert plans into suitable orders and allocate/disseminate orders to implementers.

At the "Management" level, orders received from the Executive level could be immediately converted into orders specific to the Employee level, communications channels could be selected for covertness and/or intervisibility, sequentiality of transmittals could be established and orders transmitted.

It would of course be necessary (Or would it?) to somehow explain to those remaining "Executives" (personnel requirements would be diminished) what future events have been predicted, their probabilities, what the recommended course of action is, its probability of success, and the recommended timing for transmittal of orders to assure maintenance of maximum success probability. It might also be desirable for the system to identify additional data needs required to improve probabilistic predictions.

Similarly, at the Management level, orders and their rationales might have to be explained to the remaining human participants.

The parenthetical question in the above paragraph refers to an issue which is being dealt with in today's aircraft cockpit. It is becoming accepted practice today that nothing potentially requiring human judgment be allowed to occur automatically, or without pilot concurrence. This position is reflective of our current lack of confidence in cognitive-decision aiding systems and conservatism in the application of new technology to a traditionally human role. Such hesitancy has characterized most proposed innovations to aviation crew station information management such as the introduction of multi-function displays instead of "steam gauges" and the use of three dimensional display formats. Hesitancy has also no doubt attached to the adoption of newer technology in other applications as well. This too shall pass.

These potential applications of newer data management technology are not unlike what are being considered for combat aviators. Information being considered and evaluated for implementation for aviators consists of recommended flight routes, automated formulation of attack plans, recommended observation and attack positions and timing and recommended communications channels. These systems can take into account threat data, threat and self position as well as the activities of cooperative forces and successes or failures of enemy force projections. Data is automatically selected for increased display prominence, or "conspicuity", to the crew to assure its timely perception.

For the Tank crew member or commander, information could consist of probabilistic data on predicted enemy location and the success probability of one-on-one or combined forces attack. Platform/weapon selection as well as a recommended in-range intercept course and firing position could also be provided based upon somewhat aged sensor data on previous location and probabilistic vector of enemy movement, terrain characteristics, local weather, battlefield obscurants, known placement of mines and mobility characteristics of the firing platform(s).

## 4. PROBLEMS AND THEIR RESOLUTION

The above recommendations are only limited by our imagination and our willingness to entertain the possibility that our machines might be able to accomplish tasks traditionally undertaken by humans.

Limitations of data on events transpiring in the battlefield are rapidly diminishing. The danger to be avoided is being overcome with data.

But the achievement of useful systems for the battlefield is not without problems-problems quite similar to those we are facing in providing cognitive decision assisting systems for the aircraft operator.

A vexing problem, elsewhere addressed during this workshop, is: Which responsibilities do we assign to men and which do we assign to the machine system?

This paper can only contribute to answering that question by harking back to the human limitations (and approaches to their resolution) mentioned earlier. By restating the question in terms of which tasks require interpretation of quantities of data, discernment of probabilistic unifying characteristics within the data, dispassionate prediction of future relevant events and the development of information to guide action, we have the beginnings of a set of criteria which might be useful in task allocation. Those which meet these criteria should obviously be considered for allocation to machines. Those which don't are candidates for man. Others in this field would take the opposite approach. We disagree with the others.

Another problem of concern to this meeting is whether there are some types of tasks which should never be entrusted to the Human-Electronic Team. The position taken by the authors of this paper is that any task meeting the above criteria should initially be entrusted to the team, with a larger proportion of them assigned to the electronic crew member, or machines. The logical alternatives include all tasks be assigned to man and man and machine share task responsibilities. By not taking the initial position that all tasks be assigned to machines, we may be flagrantly displaying the limitation in our human ability to accept risk. Until it is convincingly demonstrated that the success probabilities predicted by our decision-aiding systems are unrealistic, our all too human tendencies should be recognized and resisted.

In answer to the question of whether current development activities address the teaming

43

issues, we must say only partially. Recent evidence indicates that cognitive decision-aiding systems are being examined for application to rotor craft, tanks, and remotely operated vehicles. Applications of this technology to the upper command levels, where they are also potentially of great value, appear to be absent from current planning. It is also not apparent that any of these potential applications involve the application of fuzzy logic or are concerned with improvements in risk acceptance.

Another serious problem which exists but has not been clearly stated as a concern of this conference is the question of how to develop expert systems through the process of knowledge acquisition when knowledge in the application of the systems projected for use on the battlefield does not reside in any experts today? Whom do we query for expert advice on how the systems should perform? This is a question which has been asked at the highest levels without satisfactory response.

It is our belief that, as with any new technology, the answer to this question will come from experience in its use. Thus, its full potential will not be realized until experience has been gained with its capabilities. This does not mean that imaginative suggestions should be dismissed because they have not been proven. Rather, suggestions should be readily entertained in the laboratory or in the simulated environment where they can be thoroughly examined for their worth. Doing otherwise would be yet another example in which man is dominated by risk-reduction mechanisms.

## 5. CONCLUSIONS

The purpose of this gathering is to consider several questions. Among them are what roles should be assigned to man and what roles should be assigned to machines; as well as how to establish criteria for role selection and assignment. While that purpose is laudable for the aviation context, it is perhaps too limited. We have argued that the problems of decision-making in aircraft applications are not significantly dissimilar from decision-making problems in all other levels of battlefield information management. To help in answering these questions we have offered a conceptual framework and guidelines. We question the advisability of limiting the application of this emerging technology to aviators and tank crewmen and suggest that opportunities to exploit its possible contributions be extended to all combat operations in which human decision-making capabilities are employed.

Thus we are recommending that the role of the combat aviator is not the only place that advanced technology encompassing fuzzy logic

can and should be applied. Such technology may be useful in improving "cognitive decision making" for all players and at all levels in the modern electronic battlefield.

We are also visualizing an unstated, but logically inescapable, possibility that has been entertained by writers of science fiction. That possibility is that if the technology of data management, decision making and the generation of information (as defined in this paper) progresses to one of its many logical conclusions, the time may come when man will be absent from the battlefield. That battlefield may be controlled and conducted by machines; only machines will be destroyed during warfare and man will discover which side won or lost when the machine announces the outcome and tells him what to do next.

## 6. REFERENCES

1. Edwards, W., "Conservatism in Human Information Processing." in Kleinmuntz, B. (1968). (ed.): "Formal Representation of Human Judgment", New York: Wiley, pp.17-52.

2. Zadeh, L.A. "Fuzzy sets. "Information and Control", 6, 1965, pp. 338-353.

3. Bhatnagar, R. & Kanal, L.N., "Models of Enquiry and Formalisms for Approximate Reasoning", in Zadeh, L.A. & Kacprzyk, J. .(eds.): "Fuzzy Logic for the Management of Uncertainty," New York: Wiley, 1992 (ISBN 0-471-54799-9), pp. 29-54.

4. Zadeh, L.A. & Kacprzyk, J. .(eds.): "Fuzzy Logic for the Management of Uncertainty." New York: Wiley. 1992 (ISBN 0-471-54799-9)

# SESSION II - KNOWLEDGE ENGINEERING METHODOLOGY

## SYNOPSIS

The papers in this important technical section are concerned with the methodology and technology of intelligent, knowledge-based systems. They describe the capture, and the engineering, of the knowledge that is necessary to make EC an intelligent assistant and advisor. The methodologies and technical issues covered include: logical inferencing for task automation initiation (Paper 7); case-based reasoning for complex system diagnostics (Paper 8); and procedures for knowledge acquisition (Papers 9 & 10). *Paper 7* proposes an expert system for real-time dynamic HE-C task allocation, and argues that the system accomplishes the task in a manner that builds trust. Condition (or state) information regarding the aircraft, crew and surrounding environment, based on domain expert knowledge, is used to determine the task initiation. The knowledge acquisition (KA) methods are described for capturing, using a wordprocessor, the decisions and the action criteria for a rule-based expert system. A formula is used to predict the effect on trust, over time, for different fault sizes. It is argued, that since the knowledge bases do not change during operation, the task initiations are predictable, and hence trustworthy. Trust is also engendered by the parallel nature of the implementing discrimination network, which operates speedily, in real-time; by the flexibility ensured by inferencing paths that deconflict automation tasks and ensure on-time execution; and by task scheduling that requires crew consent for initiating and continuing tasks. *Paper 8* discusses the acceptance of knowledge-based systems (KBS) approaches to airline aircraft maintenance. It summarises lessons learnt from a first generation technology demonstrator system for engineering diagnostic support. A convincing argument is advanced for the use of the case-based reasoning (CBR) technique in support of complex system diagnostics that need to be made under time constraints. CBR overcomes criticisms of earlier rule-based, and model-based AI approaches, by substantially reducing the intrusiveness of the knowledge engineering input required. CBR adapts solutions that were used to solve old problems, and draws upon understanding of human semantic and episodic memory. This provides an intuitive feel to the technology, which complements human thinking, reasoning and problem solving. *Paper 9* focuses on the tools and techniques required to design real-time airborne tactical Decision Support Systems (DSS), based on KBS technology. This paper reports the development of tactical advisors for use by the commander in future Anti-Submarine Warfare (ASW) and Anti-Surface Warfare (ASuW) aircraft. ASW/ASuW DSS provide aiding, whilst ensuring that the human operator retains control, in circumstances where total reliance on automation is either inappropriate or impossible. The authors describe a distributed version of the MUSE real-time KBS tool-kit (DMUSE). They propose a top-level, generic structure for a real-time KBS tactical decision aid, that takes advantage of the multi-process structure of DMUSE. A KA automated toolkit, based on KADS, is described for supporting this real-time KBS structure. The quality of aiding of situation assessment has proven to be a key operational factor. Improving KA for aiding situation assessment has required the development of KA interview techniques based on Repertory Grid, Principal Components Analysis, and Case-Based Reasoning methodologies. The Critical Decision Method (CDM), based on recognition-primed decision theory, is considered to be particularly sensitive to acquiring the implicit information underlying tactical decisions. *Paper 10* discusses two fundamental problems in the KA process, which are critical to the design of DSS. The problems are, making the best use of contact time with experts, and the lack of direct expertise in the operation of future, non-mature systems. A KA methodology (REKAP) is described which offers solutions through the use of structured analysis and pre-defined models of problem solving inferences. The methodology provides a conceptual model of the knowledge possessed by domain experts that is refined during acquisition to meet the specific domain needs. REKAP, and the tools which support it, are described with reference to an ASuW demonstration application. REKAP provides a framework to visualise the acquired knowledge during acquisition, which helps consistency checking, and improves the productivity of KA sessions. In general, it seems that, whilst problems of problem solving inferencing and real-time processing continue to provide substantive technical issues, the focus has shifted towards  improving DSS useability, user confidence, and system effectiveness by using KA methods. Knowledge engineering techniques, such as CBR, CDM, and REKAP, help capture and exploit expert knowledge and strategies for dealing with uncertainty. They generate more "believable" DSS, that support the operator's implicit knowledge and tacit skills, in ways which the user can easily understand, and in a manner which is more likely to engender trust in DSS advice.

# TRUST-ENHANCING SENSOR AND INFORMATION FUSION
# FOR KNOWLEDGE-BASED COCKPIT DECISION AIDS

Peter G. Raeth
Wright Laboratory, WL/FIPA
Pilot/Vehicle Interface Technology Branch
2210 Eighth St, Ste 1
Wright-Patterson AFB Ohio 45433-7511 USA
raethpg@wl.wpafb.af.mil

James L. Noyes
Wittenberg University
Mathematics and Computer Science Department
Springfield Ohio 45501-0720 USA
noyes@wittenberg.edu

Anthony J. Montecalvo
Veda Incorporated
5200 Springfield Pike Suite 200
Dayton Ohio 45431-1255 USA

## ABSTRACT

The growth in complexity of modern combat compares unfavorably with the decreasing cost-effectiveness of additional crew. Cost-effectiveness issues cause military planners to remove crewmembers from the weapon system. However, reduced crew size means that the remaining crew must still survive to perform their missions successfully in combat situations that stress the mental and physical extremes of the best people. One way to resolve this dichotomy is to let the weapon system itself perform mundane chores while the human crew acts as a battle manager. (This is consistent with Sheridan [1] and Pawlowski & Mitchell [2].) At variance with this concept is the lack of trust humans tend to have for embedded automation. This lack of trust is typically based on several misperceptions about intelligent systems: 1) their unpredictability, 2) their inflexibility and slowness in realistic combat scenarios, and 3) the crew's loss of control over the weapon system when the automation is active.

This paper proposes a framework for initiating tasks to be performed by intelligent systems based on raw data available in the weapon system. This data comes from sensors, instruments, or off-board platforms. Time-dependent raw data is used to determine the conditions evident in the weapon system, crew, and combat situation. Condition information is used as ground truth to determine the tasks to be initiated, either automatically or with crew consent. The automation is "programmed" as a knowledge-base developed by operational, mission, human factors, and weapon system experts. The nature of the automation does not change during operation. This ensures predictability. Speed is ensured by the underlying scalability of the inferencing techniques. Crew control is ensured by a task scheduling method that requires human consent for initiating and continuing tasks. Flexibility is ensured by inferencing paths that deconflict automation tasks and ensure on-time execution.

## BASIC EXPERT SYSTEMS FORMULATION

The value-of expert systems to help solve a variety of diagnostic and advisory needs has been well-demonstrated over the last two decades. Sometimes, a large number of rules must be continuously checked in real-time due to stringent requirements imposed by the problem. Because of this, the expert system must scale so that the timing demands can be met in spite of expanded knowledge volume. This paper presents expert systems techniques that meet this demand by scaling in a natural way for parallel processing architectures.

Techniques for two expert systems will be introduced. Domain experts using simple wordprocessors can generate the knowledge bases for these expert systems. The result is automatically translated into appropriate discrimination networks and matrices for use by the inference engines.

CONDITION ANALYSIS: observes raw data gathered over time from various sources to determine specific conditions evident in the aircraft, crew, and surrounding environment.

ACTION DETERMINATION: selects actions based on the conditions evident in the aircraft, crew and surrounding environment. The conditions are decision criteria that, when satisfied, lead to action.

Besides the information in this paper, related discussion may be found in Raeth [3,4] and Noyes [5,6]. Implementation details are in Raeth, Montecalvo, & Noyes [7] and Raeth, Noyes, & Motecalvo [8].

## COLLECTING & ORGANIZING RAW DATA

There are two types of raw data needed by the condition analysis inference engine, measured and derived. Measured data come directly from sampling various data sources. Each data source becomes a data matrix column. Derived data are the result of computations performed on measured data. Derived data are merged with measured data to create additional columns. The time stamps are the rows. The conditions are determined by observing the raw data available over time in the cockpit.

## CONDITIONS FROM RAW DATA

Two things must happen to support the condition analysis phase. First, the knowledge of domain experts must be captured. Second, this knowledge must be used consistently, with the raw data, to determine when conditions are active.

Based on extensive interviews with experienced domain experts, three basic categories of information

make up this knowledge base:

CONDITIONS: internal/external aspects of the aircraft, mission, and human situation that are active for a given time.

EVENTS: raw data correlations that must be observed in time sequence before a condition is determined to be active.

SPECIFICATION SETS: triplets that identify a data source, its expected value, and a numerical comparison operator (entered as $<$, $>$, $=$, $<=$, $=>$, or $<>$). A list of these sets makes up an event. The sets in an event must all be recognized simultaneously before the event is determined to have occurred. Specification sets are used to identify events, and events are used to identify conditions.

Besides the three categories mentioned above, two categories of related information that also must be collected:

DWELL TIME: the minimum length of time an event must last.

TRANSITION TIME: the maximum time between the start of one event and the beginning of the next event.

It is very important to devise a mechanism oriented to the domain expert for the capture of these five categories of information. The mechanism chosen must use the language of the expert and be simple enough for that person to modify the knowledge base incrementally. The authors' personal experience is that a common spreadsheet does nicely when set up on a wordprocessor as illustrated in Figure 1. The inference engine's preprocessor transforms the spreadsheet shown in Figure 1 into a discrimination network. Notice that each row of the spreadsheet represents an event with that event's specification sets. The condition named at the beginning of the row is associated with the event on that row. If there is no condition name and the line is not totally blank, then the event on that row is the next sequential event for the previously named condition. Separate linked nodes are created in the computer's memory to associate each condition with its events and each event with its specification sets.

Figure 2 shows the basic structure of the discrimination network developed from Figure 1. Physically, this network is implemented as a series of connected linked lists. The conditions are listed vertically, the events for each condition are listed horizontally, and the specification sets for each event are listed diagonally. There is no logical limit to the number of conditions, events, or specification sets. Each condition is evaluated independently. Therefore,

this framework should scale well across multiple processors as more conditions are added to the knowledge base. The network shown in Figure 2 is traversed in a forward-chaining, depth-first fashion. Each data sample is compared to the specification sets of the current event of each condition on the active search list. Once a specification set is found that does not compare as it should, the comparison process for that condition stops and the comparisons for the next condition start. If all the specification sets compare correctly and their event's dwell and transition time factors are within range, the current event is passed and the next event is made the current event for comparison purposes on the next sample. Once all events for a particular condition are passed, that condition is identified and placed on the active condition list. Active conditions remain so until one of their anti-conditions is found.

BUILDING A KNOWLEDGE BASE OF DECISIONS AND THEIR CRITERIA  The bounding of the knowledge domain and the acquisition of knowledge are two classical bottlenecks in the development of expert systems. The knowledge acquisition method discussed here is designed to deal with these two bottlenecks.

DOMAIN BOUNDING: Domain bounding is implied in the method of capturing the knowledge. A distinct boundary is drawn around each decision that could be made. This boundary encompasses the criteria corresponding to each decision captured in the knowledge base. Thus, the "domain boundary" (as understood in the classical sense) is actually composed of many minute boundaries. Later, you will see how these many small boundaries are kept from conflicting with each other, although some of them overlap.

KNOWLEDGE ACQUISITION:  Two kinds of knowledge are captured. The first is composed of the decisions and their resulting tasks. The second is each decision's corresponding criteria. Overlapping criteria are resolved automatically by the inference engine's preprocessor. Decision criteria are derived from the existence of conditions. A list of criteria must be satisfied before its related recommendation or action is undertaken. A mechanism for capturing and organizing this type of knowledge can be constructed with a wordprocessor. There are several elements to a decision description:

- DECISION NAME

- ON TASKS: list of tasks to be performed when the decision is made

- OFF TASKS: list of tasks to be performed when the decision is no longer appropriate

47

- CRITERIA LIST: what criteria must be met before the decision is appropriate, these are the conditions that must be active before the rule can be fired

## FORMULATING A DECISION/CRITERIA EXPERT SYSTEM

The formulation for the evaluation of decision/criteria knowledge bases depends upon a criteria vector, a response (action) vector, and a set of relationships between criteria and responses.

The criteria vector c is a vector of m Boolean (True or False) variables. These criteria are the conditions discovered by the condition analysis expert system discussed earlier.

A set of n rules define a rule vector r, relating the criteria and response vectors, defines the on-board expert system that will advise the pilot and, with the pilot's consent, act on their behalf. Each rule can be formulated as a conjunction of simple Boolean criteria that lead to a set of actions. If all a given rule's criteria are true, a given action will result. (Note that an "action" could be composed of any number of activities.) All possible actions define an action vector a of size p. Each rule is expected to involve only a relatively small number of m possible criteria. The rule-base is built off-line, and not modified during the evaluation process.

In a typical rule-based expert system, the inference engine performs three standard operations:

1) the match operation matches the criteria against the rules to see which actions could occur

2) the resolve operation chooses which of these actions will actually occur

3) the execute operation actually generates the appropriate actions and updates working memory

Two methods for rule evaluation will now be introduced. Each rule is evaluated independently so the methods should scale well across parallel processors. If required, different levels of parallelism could be employed. If the processing time is not fast enough, then rules having the same priority could be grouped according to their number of criteria to equalize the work among the processors, as discussed by Tout and Evans [9].

## DECISION/CRITERIA EVALUATION: METHOD-1

The simplest method for this expert system evaluation assumes that the rules and their criteria are listed in priority order. This is equivalent to a priority-oriented backward chaining method. This is the obvious choice when $n \ll m$ and no other assumptions are made about available data. Because no OR-logic is present in a given rule, Method-1 stops with the first $c_i$ = False

(or first $c_i$ = True in the case of $\sim c_i$). If these rules were ranked and evaluated from highest to lowest priority, then the first action produced (if any) would be the most important from the pilot's point of view.

The rules could be represented efficiently by using three vectors: the previously discussed action vector a whose elements each point to a specific task to be completed, a query vector q, identifying which criteria have to be checked, and an index vector End, that delimits the criteria that appear in each rule. Here q employs positive integers to indicate criteria indices. Negative indices indicate criteria complements (NOT-criteria). This allows for direct and very fast access to the c vector stored on the blackboard (only one internal integer multiplication and addition are needed to compute any cell address). If parallel processors are used, this Boolean criteria vector c can be accessed from the blackboard by all processors. If multicomputers are used, c would be communicated to the local memory of each processor and this communication time needs to be considered, according to Lester [10]. Each processor uses components from the query vector q. Note that the number of elements in vector q is equal to the total number of criteria in all of the rules. Vector End has n elements, the total number of rules.

This method yields Algorithm-1.

```
Forall i := 1 to n do in parallel
  begin
    if i = 1
      then j := 1
      else j := End_{i-1} + 1;
    Fired := TRUE;
    while j ≤ End_i and Fired do
      begin
        k := q_j;
        if k ≥ 0 and not c_k
          then Fired := FALSE
          else if k < 0 and c_{-k}
          then Fired := FALSE;
        j := j + 1
      end;
    if Fired then perform action a_j
  end

Algorithm-1.  Implements Method-1
```

## DECISION/CRITERIA EVALUATION: METHOD-2

The previous method does not take advantage of searching in any informed way whenever the raw data (and hence a criterion) changes because the indexing is in the opposite direction from rule to criterion. A second, combined forward-backward chaining method, could be used to check only the rules whose criteria values have changed since the last evaluation of the

rule-base. To do this, one also could index in the opposite direction, checking only the rules having newly changed criteria. The forward phase identifies the changed criteria and rules that use these criteria. The backward phase is the same as before with presumably fewer rules to process.

Assuming certain criteria were the only ones whose truth value changed, their **NeedToCheck** components would be set to True in the blackboard. **NeedToCheck** components are reset to False after their rules had been re-evaluated. The **Last** vector is similar to the **End** vector of Method-1. It's indices point to blocks of rules listed in <u>vector</u> v. The idea is for criteria to point to their rules. Then it is possible to re-evaluate only those rules for which criteria truth values have changed. Once these rules are identified, the actual criteria checking occurs as in Algorithm-1.

<u>DEALING WITH UNCERTAINTY</u> In practice, one or more sensor failures may lead to undetermined (uncertain) components of the raw data source vector s, which may lead to one or more unknown truth values in the criteria vector c. For every rule, one of three situations must hold at time-step $t_k$:

1) the truth value of all its criteria are known

2) there are criteria with unknown truth values, but at least one of the known criteria fails to be satisfied

3) all of the known criteria are satisfied, but there are still criteria of unknown truth value

The first two situations are easily addressed, since it can be exactly determined whether or not the rule will fire. In the first case, the rule will fire or not depending on the truth value of its criteria. In the second case it will not fire. In the third situation, criteria with unknown truth values determine whether the rule will fire or not. It is possible to <u>report</u> a <u>possible action</u> by simply keeping count of the number of criteria that are unknown for the given rule. A possible action occurs if a rule's criteria are either True or Unknown. A confidence ratio may be computed by dividing the number of unkown-value criteria by the total number of criteria in the rule.

<u>IMPACT ON TRUST</u> In their study, Lee and Moray [11] derived the following equation for predicting the amount of trust a human operator will have in embedded automated assistants:

$$trust(t) = 0.570*trust(t-1) + 0.062*performance(t) - 0.062*0.210*performance(t-1) - 0.740*fault(t) + 0.740*0.400*fault(t-1)$$

where: t = time index
t-1 = previous time index
performance = % of top performance
fault = % response error (Ex: operator sets 10 rpm, system delivers 10 +/- fault rpm)
trust = % of absolute trust

Figure 3 shows this equation plotted for degraded performance when there is no observable fault. Figure 4 shows this equation plotted for a maximum 100% error gradually increasing as performance degrades. While there is no room to show other plots, the authors have noted the following impact of fault on trust, as computed via the equation:

a) fault size is the primary driver of trust loss

b) the higher the fault size, the slower trust is recovered

c) trust recovers to 100% once performance is restored and faults are eliminated

d) the faster faultsize increases, the faster trust is lost

e) trust builds quickly from an initial 0% if there are no faults and the system operates as expected

These computations appear to propose that if the knowledge-based system described here were used to monitor errors and failures and if it were to call for human or automated recovery, then the severity and impact of those errors could be minimized. Thus, if one accepts the equation, trust would be enhanced.

<u>CLOSING COMMENTS</u> One might comment that the techniques described here sound a lot like AND/OR logic that could be implemented in hardware or via traditional sequential code. This is true for specific and unchanging condition and decision cases. However, the complexity of construction in hardware and sequential code increases with the complexity of the conditions' specifications, especially conditions that are time dependent and based on combinations of samples from several sensors. With expert systems techniques, one need only fill in a spreadsheet implemented on a wordprocessor. This permits going from concept to reality very quickly. Figure 5 shows a top-level view of how the decision support capability introduced in this paper might be integrated into an embedded processing environment.

The methods discussed in this paper are predictable because the underlying knowledgebases do not change during operation and they directly reflect the flight operations manuals. Thus, they are predicated on procedures and concepts of operations with which humans are familiar. They operate in real-time due to their parallel nature. This parallel nature exists at all

```
 -----------------------------------------------------------------------------
|                              |  |EVENT  TIMING      ||EVENT SPECIFICATION SETS    |
| CONDITION NAME               | \|DWELL |TRANSITION  ||S # |VALUE|COP||S # |VALUE|COP|
|------------------------------|--|------|-----------||----|-----|---||----|-----|---|
 main & backup gen failure        100     0             3     5    >=    4     5    >=
 main & backup gen failure    \   100     0             3     2    <=
 main & backup gen failure    \   100     0             4     2    <=

 master caution blinking            0     0             2     5    >=
                                    0     0             2     2    <=
                                    0     0             2     5    >=
                                    0     0             2     2    <=
 master caution blinking      \   100     0             2     5    >=
 master caution blinking      \   100     0             2     2    <=
```

S #:      Data Source #          VALUE:     Expected Value

COP:      Comparison Operator
\:        Anti-Condition (how to tell when the named
          condition is no longer active)

Figure 1.    Knowledge capture spreadsheet for the aircraft
             condition analysis inference engine, with examples

```
                              O (any number of specification sets can be added to each event)
                             /
                           O 2nd specification set (data source #, expected value, comparison op)
                          /
                        O 1st specification set (data source #, expected value, comparison op)
                       /
1st          O-----O----------------O------ (any number of events can be added)
Condition    |     1st event        2nd event
             |     dwell time        dwell time
             |     transition time   transition time
             |
2nd Condition O
             |
             |
             O (any number of conditions can be added)
             |
```
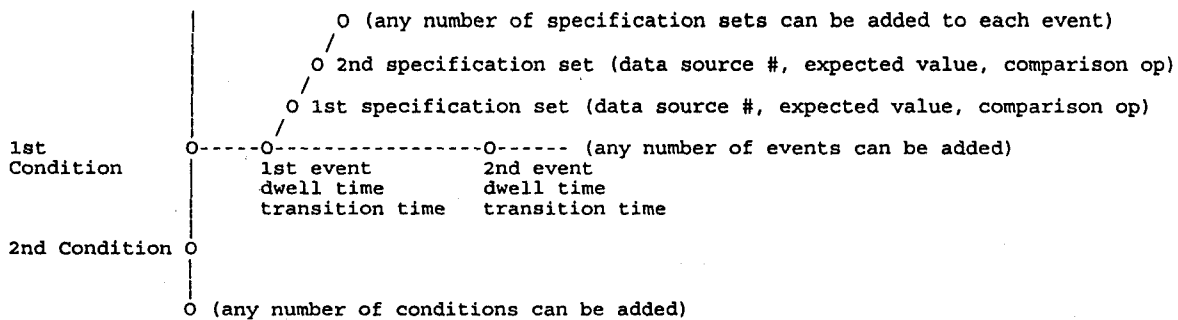
Figure 2. Structure of the discrimination network for
          the condition analysis inference engine
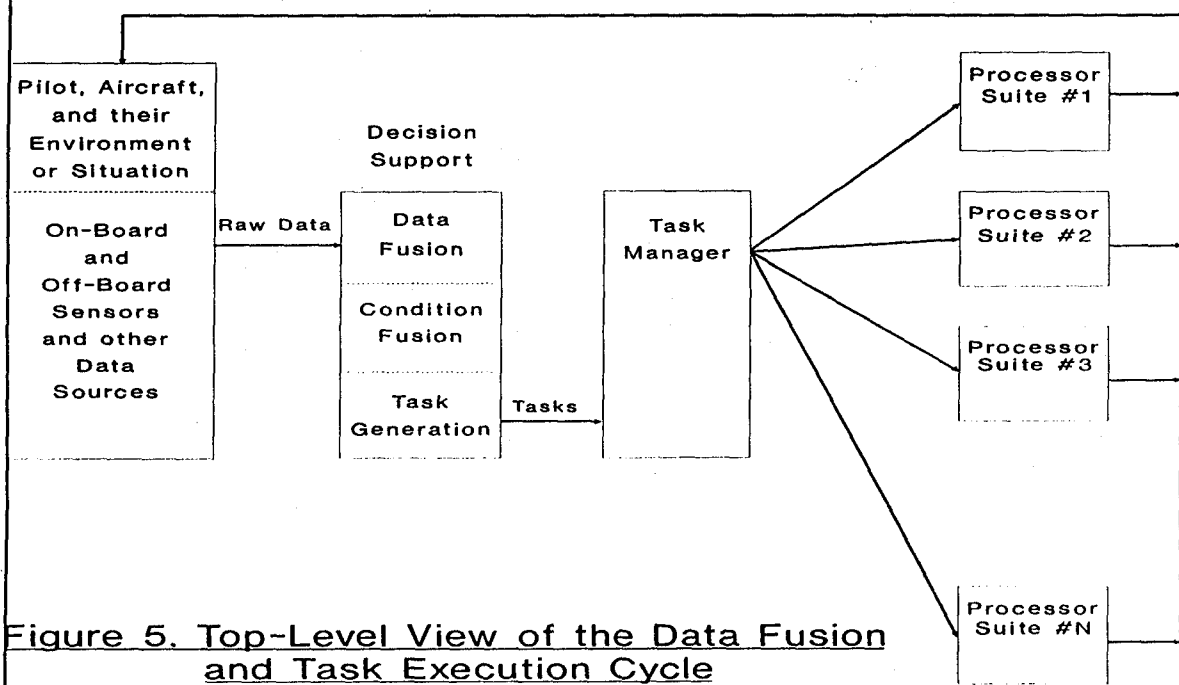```



Figure 5. Top-Level View of the Data Fusion
          and Task Execution Cycle

levels of the design. The methods are flexible because tasks are spawned, deconflicted, and scheduled as needed during operation for on-time execution. They make no assumptions about the sequence of activities in the observed environment. The crew retains control because they can negate a task at any time and because the rule-base supports pre-mission and during-mission tailoring. (Any rule can contain criteria for crew-permission given pre- and during-mission.) One use of the ideas presented in this paper is that if the pilot fails to react to traditional cockpit cautionary signals concerning failure modes, the automation will begin to react, giving additional warnings and suggestions. In the long run, this reaction could become a dynamic allocation of the remedy procedure from pilot to computer.

## Development of Trust Over Time
### (Maximum Fault Size = 0%)



Figure 3

## Development of Trust Over Time
### (Maximum Fault Size = 100%)



Figure 4

As Harvey [12] observed after the V-22 Osprey Tiltrotor crash, the operator must have timely and pertinent information to be aware of, and recover from, system malfunctions. If the operator cannot cope with the situation, the aircraft computers need sufficient tasking capability to take charge. This paper has introduced a scalable expert systems framework for accomplishing that end in a manner that builds trust.

BIBLIOGRAPHY

1. Sheridan, T.B. "Supervisory Control" in G. Salvendy (ed) Handbook of Human Factors, New York, NY: Wiley, 1988

2. Pawlowski, Thomas J. and Christine M. Mitchell "Direct Manipulation to Facilitate Supervisory Control and Intent Inferencing in Complex Dynamic Systems", Proc: IEEE Int Conf on Systems, Man, and Cybernetics, Oct 91

3. Raeth, Peter G. (editor and contributing author) Expert Systems: A Software Methodology for Modern Applications Washington, DC; IEEE Computer Society Press; 1990

4. Raeth, Peter G. "Expert Systems in Process Observation and Control", AI Expert, Sep/Dec 90

5. Noyes, James L. Expert System Rule-Base Evaluation Using Real-Time Parallel Processing Alexandria, VA: National Technical Information Service, Report # WL-TR-93-3098, 1993

6. Noyes, James L. Artificial Intelligence With Common Lisp: Fundamentals of Symbolic and Numeric Processing Lexington, MA: DC Heath, 1991

7. Raeth, Peter G; Anthony J. Montecalvo; and James L. Noyes Scalable Expert Systems for Adding Crisp Knowledge to Pilot Vehicle Interfaces, Cameron Station, VA: National Technical Information Service, Report # WL-TR-??-????, to be published

8. Raeth, Peter G; James L. Noyes; Anthony J. Montecalvo "A Scalable Framework for Adding Crisp Knowledge to Pilot/Vehicle Interfaces", Proc: IEEE Conf on Man, Systems, and Cybernetics, Oct 94

9. Tout, K. R. and D. J. Evans, "Parallel Forward Chaining Technique with Dynamic Scheduling, for Rule-Based Expert Systems," Parallel Computing, Vol 18, #8, pp 913-930, Aug 1992

10. Lester, Bruce P., The Art of Parallel Programming, Prentice Hall, Englewood Cliffs, NJ, 1993

11. Lee, John and Veville Moray "Trust, Control Strategies, and Allocation of Function in Human-Machine Systems", Ergonomics, Vol 35, # 10, pp 1243-1270, Oct 92

12. Harvey, David S. "V-22 Crash Prompts Questions About Cockpit Data & Displays", Avionics, Aug 93

# CASE-BASED REASONING AND AIRCRAFT SYSTEMS TROUBLESHOOTING: NEW SOLUTIONS FROM OLD

## R V Magaldi

Information Management Research, AI Unit, British Airways Plc, Heathrow Airport, England

### Summary

Knowledge-based approaches to the solution of aircraft maintenance problems have begun to gain acceptance within the airline business. The industry began several years ago by using Artificial Intelligence (AI) to address complex maintenance and scheduling problems. More recently, focus has shifted to other important areas of engineering, such as system condition monitoring and failure diagnostics. However, questions have been raised as to the degree of trust that can be, or should be placed on the solutions and advice given by such systems, particularly in an arena where complex legislative, social, professional and ethical factors have to be accounted for.

Case-based reasoning (CBR) has recently emerged as a good candidate second generation technology, both to support complex system diagnostics, particularly when aiding human decision making within time constrained operations and to overcome some of the earlier criticisms made against AI based approaches. However, a number of human factor and technical issues still remain to be resolved before the apparent potential of CBR technology can be both accepted, and exploited within the airline industry. This paper emphasises some of the findings and lessons learned during CBR based work at British Airways, and also describes some of the authors emerging views of how such systems must be planned and developed in the future.

**Keywords**: semantic memory, episodic memory, learning, case-based reasoning.

**Introduction**. The complex and dynamic environment in which aircraft system fault finding takes place should provide an ideal area for the application of AI technology, particularly for the employment of knowledge-based systems. The operational basis of these resting on the human expert knowledge contained within them, and the manipulation of this to solve problems. Systems employing such techniques have already demonstrated their potential within other areas of industry and commerce where significant claims of increased operational efficiency and cost saving have been made.

Given that these benefits are transferable, then there exists great potential for supporting and increasing profitable aircraft operations within the airline industry, whilst maintaining high standards of passenger service and operational safety. Knowledge-based or expert systems should as key drivers of this process, be well positioned to provide a wide range of intelligent decision support facilities to aircraft maintenance staff. However, the reality is that operational deployment such systems within airline

engineering, have with certain exceptions been confined to maintenance planning and scheduling tasks [1] [2].

If AI based engineering support systems are showing growth and usage within other sectors [3], why have they received little attention or acceptance from within the airline engineering community?, when there seem to be such clear organisational and business advantages to be gained through their use. The phenomenon is a complex one, and contains as many psychological, organisational and legalistic components as it does technical ones. Such factors can include:

o Reluctance of individuals to share knowledge and experience.

o View that maintenance manuals should be the only authoritative source for diagnostic procedures.

o Fear that knowledge-based systems could make recommendations that run counter too authorised working practices.

o Problems of system ownership, and responsibilities for system maintenance.

o Accountability, who takes the blame for errors and failures in system output ?

This list is not exhaustive, but serves to highlight some of the major concerns that airline maintenance staff commonly express.

It is clear from these examples that the various AI technologies that we might wish to propose as solutions to particular maintenance engineering requirements, have more than just a technical impact. Clearly, there are organisational and procedural factors that must be accounted for in any proposals made. Technical solutions being fielded must harmonise with both human and organisational expectations, or be rejected.

**AI for Engineering at British Airways**. Until quite recently, British Airways has been investigating and evaluating a broad spectrum of potential AI technology application areas [4], including avionics, engine condition [8] and aircraft performance monitoring. Fig 1. Particular emphasis was made during this period, of isolating those areas of a high potential return on investment coupled with a similar likelihood of technology acceptance by the targeted user community. General organisational advantages to be

gained through the employment of AI based solutions were identified and communicated throughout the engineering community. These included :

o  Knowledge as an engineering-wide resource.

o  Wider availability of expert knowledge to less experienced staff.

o  Less vulnerability of the company to fluctuations in expert availability.

These concepts were widely supported by engineering management, and have been at the heart of any AI technology initiatives to date. Initial study findings indicated great potential existed for the application of AI technology to aircraft system and component failure diagnostics.

Various programmes have since been constructed since then to investigate the utility of particular AI approaches as aids to component fault finding and repair. Three diagnostic paradigms were considered and evaluated during this time. These were as follows :

o  Rule-based

o  Model-based

o  Case-based

These alternatives were found to possess various strengths as well as weaknesses, many of which were technology based. For example, rule-based systems were recognised as being useful in situations where lots of reliable diagnostic information was available, and they offered a low cost approach to many engineering requirements. They were found to be good at representing the semantic content of a domain, but were poor at handling episodic features.

In psychological terms, semantic memory can be described as containing those facts that we know about the world, with these arranged in some hierarchical manner. These instances, although complied from experience do not however, have much to say about our actual experiences of the world.

Life experiences, or scenarios are encoded in episodic memory. These consist of many facts or instances, which form relationships according to their co-occurrence in the same life episode, the use of facts varying according to particular stuational contexts. The inability of basic rule-based systems to make efficient use of contextual information was found to limit their potential as good decision support platforms within contextually rich decision making environments.

Model-based approaches seemed to offer deeper ways of representing domain knowledge, and because of this, gave a much richer environment for reasoning and problem solving. However, it was quickly recognised that a large investment of both time and money would be necessary to model

candidate components or systems. Coupled with this was the uncertainty in many instances, of what system features to model, and to what level. For example, in the case of a jet engine, do you use a gas-path model, a vibrational model or a structural materials model ? Obviously, all of these could be used, either individually or in combination. However, such models take time to develop and validate, although may have the potential of high eventual return on investment when completed.

Case-based reasoning, although fairly new as an established AI technology seemed to show promise as a candidate environment for the rapid development of engineering operations support systems. The technology seemed capable of handling fuzzy queries, displaying rudimentary learning characteristics and having available, rich data indexing and search mechanisms. The potential to build realisable engineering decision support systems within short time scales was quickly appreciated, not just by AI development staff but also by engineering customers ! There were a number of reasons for this:

o  Technology was commercially available off-the-shelf.

o  Software was well supported, and first class consultancy available.

o  Minimal knowledge engineering input required from engineering staff(when compared with other AI approaches).

o  No data/information/knowledge is thrown away.

o  Case-based reasoning has an intuitive appeal to most maintenance engineers, i.e., "have I seen this problem before ?".

o  Data/information/knowledge held in a case-base changes in step with the real world.

Principal of all of the above was the minimal intrusion expected during knowledge engineering stage, as already suggested maintenance engineers seldom have the time to devote to extensive knowledge engineering exercises. Access to experts is often a major difficulty when attempting systems development work in operational situations, and is made impossible by the lack of commitment to things outside established operational goals [5].

Time and opportunity for development are rare commodities in dynamic business environments, particularly within the airline business. Windows of opportunity have to be exploited without delay, whilst maintaining high standards of customer service and operational safety. Things that get in the way of these commitments are given little support.

For these reasons it was decided to move forward with case-based technology as a possible platform for future engineering decision support, particularly in the areas of system, and component failure analysis rule and model-based approaches

53

being reserved for niche applications.

**Remembering from Experience**. When dealing with the real world we rarely have the opportunity to reason in any theoretical sense about the various possible ways that we could behave within given situations. Instead we short circuit this process, and instead rely on our previous experiences of the world to guide us. In this sense, we both adopt and adapt what we have previously learned, and refer to this knowledge when dealing with new circumstances.

For example, airline operating schedules require that swift and efficient problem rectification takes place after aircraft defects are reported, if expensive delays are to be avoided. In such cases, the maintenance engineer is very likely in addition to the use of maintenance manuals, to apply deep insights gained from experience, and through this, gauge the utility of certain diagnostic approaches, or assess the merits of particular maintenance actions. That person is likely to ask questions such as "have I seen this problem before ?", and "if I have, what did I do about it ?".

However, the balance of this type of experience combined with good diagnostic ability within the engineering community can be extremely variable due to a variety of factors. These include, staff inexperience, the introduction of new technology, shift rotations, holidays, etc. It would be extremely helpful under these circumstances to be able to have a repository of positive experiences, good engineering fixes to problems, or recovery procedures. If a particular technology such as CBR is able to support these requirements, then aircraft can be efficiently maintained, and departure delays due to engineering problems minimised.

**Case-Based Reasoning: What is it?**. CBR has a number of inspirational sources, including Machine Learning and Natural Language research. Much of the theoretical work being laid down and carried out by R. Schank, who's pioneering work on memory and learning has given rise to much of the vocabulary, and theoretical underpinning's of the subject [5][6]. A good definition of case-based reasoning is as follows:

" A Case-based reasoner solves new problems by adapting solutions that were used to solve old problems" C Reisbeck & R. Schank [6]

As previously mentioned, human problem solvers often rely on previous experience to guide thinking and reasoning when encountering new or unfamiliar situations. The problem-solver recalls previously stored cases (events) and decides how similar or dissimilar these are to some current problem. If the previous case, or cases provide any insight then the problem solver tries to solve the present case using strategies known to have been effective in the past.

On the other hand, if in solving the new problem, the problem-solver finds that past cases are different then learning will take place by a process of adapting features of the old cases, and using these, construct a new case for problem solving. If the newly constructed case succeeds, then it is stored for future reference.

From a computational point of view, CBR refers to a number of interacting concepts and techniques (e.g., statistical pattern recognition [Nearest Neighbour NN, Classification and Regression Trees CART] algorithms, and data structures) that can be used to record, sort, index and retrieve cases. In addition, repair procedures are available for case modification (equivalent of hypothesis testing).

In summary, modern CBR theory is derived from research and understanding of certain human memory processes, and is based on the following assumptions:

o  Memory is episodic

o  Memory is dynamic

o  Memory is richly indexed

o  Learning is triggered by failure

o  Experience guides reasoning

These memory features give way to higher level memory strategies that are used to deal with the various situational complexities encountered in the world. Strategies to do with:

o  'Lazy' model of the brain - Don't do fundamental problem solving unless it is really necessary.

o  Reminding - "Have I seen this feature before ?".

o  Reusing old solutions, perhaps with modification.

Taking the above into account then we can formalise CBR problem handling as a step by step process, with certain features being common to any CBR system, regardless of the technology base. Particular CBR technologies use different algorithms and types of data structure, but have similar ends in mind. The essence of CBR can be expressed in the following steps:

o  Assess problem and locate key features.

o  Use key features to see if we can remember any similar situations.

o  Take the solution that was successful last time and try and re-apply it.

o  If it works, fine. Take no further action.

o  If procedure fails, modify it and try again.

o  When a modified solution works, store it as a new case.

o  After several experiences, generalise the case.

**How to Find a Case**. Various strategies can be adopted to index and through this retrieve cases. The following are fairly typical of those found in commercial CBR tools:

o Nearest Neighbour or NN - Similarity scoring using weights.

o Inductive - Automatically generated binary-tree indexing.

o Knowledge-Based - Causal model to guide binary-tree generation.

o Template - Discrimination network of database query-like filters.

**CBR System Development**. The steps concerned with CBR development are reasonably straight forward. The largest overhead being that of the initial problem analyses, coupled with case construction and data entry. Initial case indexing being dependent on factors such as the number of cases, and the complexity of retrievals required. Tuning of the case-base is an iterative procedure that continues until satisfactory levels of matching and retrieval performances are gained. Generally, the more cases that are available the greater degree of accuracy that will be achieved on any retrieval request. As more cases become available storage and retrieval strategies can be switched, say from a nearest neighbour approach to induction, if available in the particular CBR tool being used. Adaption or repair facilities also allow for the testing and storage of new cases.

**Comparison with Other Technologies**. Case-based systems are often compared with relational database technology, and this often gives rise to suggestions about their similarities, and these give rise to comments such as " the technologies seem to be doing the same thing", or "I can do that with my relational database". These comments rest on particular features of the two technologies that are similar, but only to a point. The following highlight some of those differences:

o Case-Based Systems:

- Support to fuzzy queries

- Rich indexing, support to repair rules

o Relational Databases:

- No support for fuzzy queries

- No support for repair rules

- No induction based indexing

**When to Consider CBR**. Clearly CBR isn't a panacea for all of the ills of AI technology within the airline maintenance arena, there are situations when other technologies are going to be more attractive as solution generators. Often several of these linked, are going to provide the basis for a robust solution. Experience to date has shown that CBR is more

clearly fitted to supporting certain engineering situations than others. Features that can guide assessment of CBR suitability are as follows:

o Where you have a lot of historical data.

o Where you need to extract order from complex data.

o Where experts talk in examples.

o Where there are lots of exceptions to rules.

o Where the organisation needs to deliver consistent decisions.

o Where experience is as valuable as text book knowledge.

**CBR Technology: Conclusions**. Initial technology assessment at British Airways has shown that CBR has a set of inbuilt capabilities that complement a specific range of engineering problems. These are by nature, often more than just technical in origin, and require an awareness of many competing human, organisational and operational factors when posing solutions. Benefits that CBR can bring to this arena can be summarised as follows:

o Intuitive feel and understanding of the technology by both developers and users.

o Complements human thinking, reasoning and problem solving styles.

o Nothing is discarded ! - we just index on different parts of what we store.

**CaseLine: Case-Based Engineering Diagnostics**. Part of the British Airways case-based technology assessment programme has been concerned with CaseLine, a first generation technology demonstrator, built to explore some of the operational performance requirements of an engineering diagnostic support system. The current system has the capability to aid Boeing 747-400 technical support engineers with aircraft fault diagnosis and rectification between aircraft arrival and departure. It can advise on past defects and associated recovery and repair procedures known from previous experience to be successful.

Use of the facility does not replace the obligation of the technical support engineer to refer to maintenance manuals, or the obligation to work to authorised procedures. These being very important principles, and should be taken into account when considering any type of computerised decision support to aircraft maintenance.

The CaseLine system consists of a simple interface by which the user can input diagnostic information, and control the search for available repair and recovery information. At the core of the facility are a number of defect cases (currently about 200), these describe previous failure instances and contain details of successful recovery actions. Cases are

constructed using a number of different information fields. these having particular weight vectors which are used for case-matching during case retrieval.

Three main search modes are available:

o ATA chapter

o EICAS Message

o Reported Defect

These can be used singularly or in concert to achieve a case retrieval. Information to be used can be entered as upper or lower case, and be alphanumeric or plain text. Pure ATA searches being a simple two digit number, EICAS messages being precise, but variable length alphanumeric text, and Reported Defects being variable length word strings. Retrievals are guided by a number of internal mechanisms that contribute to a nearest neighbour classification inductive guided search for cases, or both.

It is usual for either a single or a number of cases to be retrieved depending on whether "exact", or "partial" case matching has taken place. These can then be assessed and used according to their likely advantage.

In operation, the system assists the Technical Support Engineer with the retrieval of known, but often obscure defects. These have a number of often complex underlying but inconsistent causes. CaseLine aids the engineer in advising others, as to what procedures will have the highest likelihood of success in rectifying a fault. The engineer is still obligated to use the aircraft maintenance manuals as a final authority, as well as following approved maintenance procedures, but can avoid costly delays by cutting out less productive routes to fault analysis, and fault finding.

Conclusions: Lessons Learned. Primary lessons that have been learned and assimilated during the British Airways AI for Engineering study are as follows:

o First-line users of the technology as well as a manager must champion technology acceptance.

o Individual champions of technology have a short lifetime they move to other jobs, hence a need to act quickly and demonstrate success.

o Need to recognise the background, organisational and human issues that impact on technology acceptance.

o "Trojan Horse" technologies are vital to the customer confidence building and acceptance process (e.g., Case-based reasoning) if more radical technologies/solutions are to be understood and accepted by end users.

o CBR technology has to integrate well with or via other technologies such as databases, local area networks, user interface design tools and other operational systems.

o Performance is a big issue, speed of case retrieval is vital in time-constrained situations.

Clearly there must be a continuous champion from the user community of whatever technology solution is being fielded. if that person is absent then project support will wane unless there is widespread acceptance of the proposal and aims from top to bottom of the user organisation (this is a generalisation. there are obviously exceptions). The user must also see an obvious link between what is being proposed and the impact directly on jobs and operating efficiency. Technology should not intrude, or add additional tasks, and should not force the user adopt unfamiliar or uncomfortable procedures. Finally, the technology should as far as possible reflect and support the decision making styles of the users, and organisational dynamics. Compliance in these areas will give rise to greater receptiveness and acceptance of AI technologies such as CBR, particularly within the airline maintenance industry.

### References

1. B. Yansouni, "Calling the Right Shots in Aircraft Maintenance with Artificial Intelligence". Artificial Intelligence Techniques for Improving Aircraft Maintenance Efficiency, Conference Proceedings, ISBN 0903409 88844, 21 February 1991.

2. S. Smits, D. Pracht. "MOCA-A Knowledge-Based System for Airline Maintenance Scheduling". Innovative Applications of Artificial Intelligence 3, Eds. R. Smith & C. Scott, AAAI & MIT Press, Cambridge Mass & London.

3. J. J. Popolizo. "Waging the Diagnostic Wars in the Early 1990's: Part One", Artificial Intelligence Research. New Science Associates, pp 947-956. February 1989.

4. R. V. Magaldi, "The Integration of Knowledge-Based Approaches to Aircraft Systems Management and Failure Analysis", Advances in Systems Engineering for Civil and Military Avionics, Conference Proceedings, ERA Technology Ltd, ERA Report 91-0634, 20 - 21 November 1991.

5. Capt. S. B. Sloane. United States Navy. Rtrd. "The Use of Artificial Intelligence by the United States Navy: Case Study of a Failure", pp 80—92, AI Magazine Spring 1991, AAAI Press, Menlo Park CA, USA.

6. R. C. Schank. "Dynamic Memory. A theory of reminding and learning in computers and people". Cambridge University Press, 1982.

7. C. K. Riesbeck, R. C. Schank. "Inside Case-Based Reasoning", Lawrence Earlbaum Associates, 1989

8. S. Cumming, "Neural Networks for Monitoring of Engine Condition Data", Neural Computing & Applications, pp 6-102, Springer-Verlag, London Limited 1993.

Fig.1 Future IT support to aircraft maintenance



Fig. 2 CBR : Mechanisms

Fig. 3  Caseline Architecture

# Structured and Analytical Knowledge Acquisition Methods
# For Tactical KBS Decision Aids

## R D. Ellis, R Hepworth, H Howells & Lt. R E Bickerton RN
KBS Group
Man-Machine Integration Department
Aircraft Sector
Defence Research Agency
Farnborough, GU14 6TD, UK.

## Abstract

This paper describes the research undertaken by the KBS Group within Aircraft Sector at Defence Research Agency Farnborough, UK. The focus of activity is the development of tools and techniques required to design real-time airborne tactical Decision Support Systems based on KBS technology. Automated knowledge acquisition and design for KBS using the KADS methodology is presented, together with more analytical techniques aimed towards the elicitation of experts' tactical situation assessment. A generic top-level KBS design for real-time KBS decision aids is proposed, implemented in DMUSE, a KBS shell optimised for real-time performance and distributable across multiple processors.

## 1    Introduction

The definition used in this paper of a real-time tactical decision aid is something which improves the effectiveness of an expert in making decisions and performing tasks under time-pressured conditions. Such tasks are varied, from fire-fighting, process control, command and control to airborne tactical decision aids. The technologies that can be employed to improve decision making are also varied, including altering standard operating procedures, improving training and modifying the information presented to the expert decision-maker. One technology which can offer a large improvement to the decision-making process is that of expert systems, or the more widely defined Knowledge Based Systems (KBS).

The important difference between Decision Support Systems (DSS) and conventional KBS is their emphasis on working with an expert operator, not just for them or without them. In many domains, some tasks performed by human experts cannot be emulated by software as effectively, due to the expert's skills, knowledge, experience, perception and flexibility, which by their nature are difficult to specify and encode. The role of a DSS is to aid the operator, acknowledging the limitations of KBS, whilst exploiting their properties. Such aid includes the reduction of operator workload by performing low-level tasks, improving the operator's

situation assessment, monitoring task performance as well as proposing solutions. This requires an understanding of the working practices of the operator at a cognitive level and structuring the requirements for the DSS around a deep understanding of the operator's expertise. Such an understanding can only be achieved through the extensive use of Knowledge Acquisition (KA) methods, the classic bottle-neck of KBS.

Knowledge Based Systems are an attempt to encapsulate and emulate on a computer, some aspects of the ability of a human expert to perform tasks. Expertise can be encoded in a number of ways in these systems, although they are most readily characterised by their use of rules and heuristics as knowledge representations. These systems can improve the decision-making process either by replacing or de-skilling the human operators or aiding them by proposing solutions.

However, in many domains, some tasks performed by human experts cannot be emulated by software as effectively, due to experts' skills, knowledge, experience, perception and flexibility, which by their nature are difficult to specify and encode. The solution often adopted by KBS is to assume that a system working with the operator, producing imperfect advice is better than the operator performing tasks unaided. Such a result was reached by DRA in developing a prototype Airborne Anti-Submarine Warfare (ASW) KBS advisor. KBS is employed in the role of Decision Support Systems and the applications described are tactics advisors for use by the commander (Observer/TACCO) in a future ASW/ASuW aircraft. [Howells and Bickerton, 1994]

This situation is far from ideal and experience in progressing the ASW advisor closer to an operational system, and the development of a prototype Airborne Anti-Surface Warfare (ASuW) KBS advisor has led to the conclusion that effective operational KBS advisors will need to integrate the operator more fully into the KBS's own decision-making process and to provide aid to all of the cognitive activities involved in expert operator decision-making. Decision Support Systems are considered essential as the only means of ensuring the commander retains control of the proceedings rather than a total reliance on automated decision systems.

This requirement brings together the Software Engineering disciplines of real-time software development and the psychological aspects of expert

operator decision making and modelling. Bridging the traditionally wide gap between A.I. and Psychology.

This paper describes the research undertaken by the KBS group at DRA Farnborough to improve real-time KBS decision aids, tackling the problem from three aspects; real-time KBS shell improvements, KBS application design developments and especially Knowledge Acquisition - the classical bottleneck of KBS and an increasingly important area if expert operator decision-making is to be understood and incorporated into operational KBS advisors.

## 2    KBS Components

The technologies behind a real-time KBS advisor can be split under three main headings:
- The hardware and software infrastructure on which a KBS is built.
- The structural design of the KBS application.
- The knowledge acquisition and design process involved in producing the KBS application.

One of the problems associated with real-time KBS is their comparatively slow execution speed, inherent in complex heuristic and knowledge representations like rules and frames. One way to overcome this, aside from converting systems from rule-based languages to high-level procedural languages like Ada, is to increase the number of processors used in the hardware architecture on which the KBS is based. Unfortunately KBS languages, shells or toolkits tend to be suited to single processor systems, allowing communication paths to other programs or other KBS but not directly employing the concept of a KBS distributed across several processors and networked computer platforms. This requirement formed the basis for the development of DMUSE, a distributed version of MUSE, a real-time KBS toolkit funded by the KBS group. There are a number of structural paradigms from which to design KBS. MUSE employs a blackboard architecture, providing a flexible way of organising rule structures into semi-autonomous processes, activating when relevant information becomes available. DMUSE uses the mirroring of blackboard databases, objects and object methods across a network to achieve the effect of a large single KBS.

A particular design, involving separate Situation Assessment. Planning and Plan Execution modules is currently being developed which is well suited to real-time decision support. This design provides aid to an expert operator in both situation assessment and planning activities. It has been derived from a similar structure currently implemented in a prototype airborne ASuW Advisor, incorporating 9 processors on a network of 6 computer workstations, all implemented in DMUSE.

Detailed specifications for KBS Advisors are inherently difficult to generate due to the complex nature of expert knowledge and the open-ended requirements of decision aiding. The Knowledge Acquisition process has the task of both generating the requirements for a KBS and eliciting the expert knowledge to be embodied in the KBS. This is a complex and time-consuming process, often termed the classical bottleneck of KBS, and research has been undertaken to alleviate this bottleneck by

developing automated techniques to formalise the process. Interactive real-time KBS requires a detailed understanding of expert operator decision-making under time pressure in order to function effectively and the incorporation of decision-making theories and their associated knowledge acquisition strategies should provide the depth of understanding necessary to both describe and encapsulate the information and processes involved.

## 3    Real-Time KBS Decision Aid Structure

Functional commonality exists between real-time tactical KBS decision aids for different application domains. A common top-level design for these systems has been sought, which would take advantage of DMUSE's multi-process structure, and provide an effective means of giving KBS support to operator decision-making. Figure 1 shows the main elements of a proposed generic structure for a real-time KBS tactical decision aid:



Figure 1

This structure contains two main processes; Situation Assessment (SA) and Plan Generation. Aid to an expert operator is provided by the activities of both of these processes. This is in contrast to many KBS decision aids which either concentrate on aiding operator situation assessment through automated data fusion, or aiding operator planning by suggesting viable plans. The problem with SA KBS is that the operator's own SA is coloured by their plans and goals, which may be difficult for the operator to input to the KBS and be appreciated only to a limited extent by the SA KBS. Thus a mismatch between operator and SA KBS will occur. The SA process in this structure overcomes this by having access to both the current plan and the expectations and assumptions inherent in the plan. This concept works most effectively when the KBS produces its own plans, incorporating expectations and assumptions within each plan action. Operator input to the system is by entering SA information, entering/modifying plans or by entering/modifying expectations and assumptions about the tactical situation. This design relies heavily on a detailed understanding of operator decision-making and effective knowledge acquisition.

This poses many design decisions best described by an example:

Many KBS advisors provide aid to the operator by attempting to solve the operator's tasks autonomously. These systems may either generate plans when requested by the operator or automatically when the system deems another plan necessary.

The Airborne Anti-Submarine Warfare KBS Tactical Aid (ASW TACAID) currently under development by the KBS group is an example of a planning oriented system. It is typical in allowing operators to request a plan to be generated by the KBS, which produces a sequence of plan actions. The operator is given an explanation as to why the plan is applicable and is requested to accept/reject/modify the plan. In progressing TACAID closer to an operational system it was noted that performance improvements to the planning rule-base could only come from a better representation of the tactical situation, a representation closer to that used by an expert operator. The ASW domain, as represented on a graphical tactical display is spatial and highly complex. To interpret this spatial domain in a similar way to an expert operator requires access to all the information available to the operator. The tactical display information is available to the planning rule-base, but the operator's expectations and assumptions are not.

The current ASW rule-base is structured procedurally, with rules arranged hierarchically under particular tactics, intermixing both situation assessment and planning rules. By using the proposed generic structure, splitting situation assessment away from planning, and including explicit representations of expectations and assumptions as input information to SA, together with their incorporation in generated plans, it is hoped that this KBS can be developed into an effective decision support system. However, this process relies on a detailed knowledge acquisition effort to elicit the situation assessment and planning activities of an expert ASW operator, how these can be represented in KBS form, what role in decision-making the KBS will play and how the operator will interact with the final system. Greater emphasis has been placed on knowledge acquisition for the SA component of the system, since improvements to planning can only proceed with a tactical representation nearer to that used by an expert operator.

The ASW KBS knowledge acquisition effort has benefited from the use of the earlier version of TACAID as a data gathering and interviewing tool. This has enabled the use of analytical techniques to examine operator SA within this complex spatial environment.

An initial prototype Airborne Anti-Surface Warfare KBS Tactical Aid is currently being developed, based around a design similar to that of the generic real-time KBS structure. It did not benefit from an earlier prototype, making analytical techniques difficult to apply. Instead, formal methods based on the European KADS design methodology have been used, and early indications are that knowledge elicited from the domain experts transfers well into both SA and planning processes.

# 4    Knowledge Acquisition

This section describes three different approaches to knowledge acquisition, linked by the common goal of producing a real-time KBS of the structure described earlier. Each attempts to bridge the gap between the software engineering concerns of producing a software system which aids expert operators, and the psychological concerns of how expert operators perform tasks and thus how best to support them with KBS technology.

## 4.1    Tool Supported Knowledge Acquisition using KADS KA Methodology

In order to reduce the considerable amount of time taken in the KA process for complex real-time KBS, one of the solutions was to automate the process of knowledge elicitation, assimilation and implementation into KBS software. An automated KA toolkit was chosen, called ProtoKEW, which was based on the ESPRIT funded KADS (Knowledge Acquisition and Documentation Structuring) [Wielinga et al., 1992] KBS knowledge acquisition and design methodology. This was interfaced, via a translator, into the MUSE KBS toolkit - as a target implementation environment. This method was evaluated through the design and implementation of the ASuW KBS advisor project. The KADS KBS design methodology advocates the use of 'Interpretation Models' to describe expert decision making. These models attempt to represent the processes involved in certain types of problem-solving, such as diagnosis or classification. These models guide the knowledge acquisition process to populate and instantiate the relevant interpretation model. The ASuW KBS uses two directive models, Situation Assessment and Planning. KADS also divides expert knowledge into 4 separate types or layers, further partitioning and structuring the knowledge acquisition process. These layers are:-

Domain Layer, which describes entities in the expert's environment.

Inference Layer, which describes simple inferences which can be made about entities.

Task Layer, which groups the inferences into structures which tackle the various activities performed by the expert

Strategy Layer, which determines when a switch in task ordering should be made to adapt to changes in the situation.

## 4.2    Analytical Methods for Situation Assessment KA

The formal methods used to develop the prototype ASuW KBS Advisor were sufficient to produce an initial KBS. Experience in progressing an ASW KBS Advisor from this prototype stage has shown the need for a more analytical approach to KA. The ASW domain, as presented to an expert operator, is a complex spatial environment from which time-pressured decisions are made. Expert operators' perception of tactical situations in this domain is key to determining correct plan actions, but experts find the expression of this perception in the detail necessary for a KBS implementation difficult, if not impossible. It was found that most of the deficiencies of the prototype ASW KBS Advisor were attributable to a poor representation of the tactical situation. In order to

remedy this, and to increase integration between expert operator and KBS decision aiding, analytical KA techniques were used to elicit the properties of an ASW expert's perception of tactical situations, through the performance of an easily measurable task. The following analytical techniques were used, namely Repertory Grid (for analysis more than elicitation), Machine Learning (Induction) and Principal Component Analysis.

At present, output from these analytical techniques has been used to refine the parameter measurements used to describe the ASW tactical environment. Several revisions of these parameters have been made, each followed by a number of trial sorties with expert operators. Increased understanding of the ASW operators' interpretation of tactical situations has resulted in additional methods of providing aid from the Situation Assessment activity of the ASW Advisor system, not exclusively tied to using KBS techniques. It is hoped that these analytical methods will generalise to progress other decision aids, such as the ASuW Advisor beyond the initial prototype stage.

### 4.2.1 Repertory Grid and Principal Component Analysis for Situation Assessment KA

Repertory grids are derived from personal construct theory [Kelly, 1955]. This technique is a free-form recall and rating session in which the Knowledge Engineer makes inferences about the relationships among objects or elements, and the dimensions or constructs that the expert uses when describing the objects.

It should be noted that Repertory Grid as described here has been used for multi-dimensional scaling and grouping, rather than for its primary role as an attribute elicitation and entity differentiation tool [Fransella and Bannister, 1977].

Several snapshots were taken during a simulated ASW sortie and various parameters were recorded for each sonobuoy within these snapshots, together with the operator's decision to monitor them (a reflection of their tactical importance to the operator at a particular instant). Each sonobuoy (of 2 types, L and D) is labelled by a number, which is used to identify them on a hardcopy of the tactical display visible to the operator at the time of the snapshot.

Repertory Grid uses this data to group related parameters and similar sonobuoys, displaying this information as dendrograms. Figure 2 shows the parameter dendrogram. This shows (for instance) the close relationship between the Cone of Courses - 'Inside coc?', and the monitoring decision 'Monitored'.

The dendrogram of elements is shown in Figure 3. This visualisation of the analysis was presented to the expert, together with the relevant scenario snapshot. From this, further knowledge was acquired. For example, the expert saw a strong correlation between the groups of elements (buoys) shown on the dendrograms, which corresponded to their positions relative to one another. This connection had not been apparent from the numerical data. These connections are illustrated by the groupings shown in an extract of the relevant snapshot (Figure 4).

Figure 2

Figure 3

**Extract from a Scenario Snapshot**
(Annotated to show buoy groupings from element
dendrogram)
Shows the perceptual structure within a tactical sonobuoy
barrier, highlighted by Repertory Grid



Figure 4

Principal Component Analysis is a numerical technique which takes the measured parameters, describing each sonobuoy in an N dimensional space, and collapses them down into two dimensions chosen such that the sonobuoys are widely spread within those dimensions. This has the effect of highlighting regularities within the sonobuoys, grouping some and differentiating others. The technique takes the same type of data as Repertory Grid, giving comparable results. Figure 5 shows an annotated diagram of the sonobuoys from the Repertory Grid data. Monitored sonobuoys are shown in larger text. The monitored yes/no parameter has been removed - otherwise sonobuoy monitoring status would be very easy to differentiate in the diagram!



Figure 5

Difficulties in separating monitored from not-monitored sonobuoys in the diagram are a reflection of the inadequacy of the spatial representation described by the parameters to predict expert monitoring decisions.

### 4.3    KA Interview Technique

KA technique is differentiated from KA methodology. Methodology provides organisation and structure for the process of knowledge acquisition. However, in any subjective process like KA, the skills and experience of the practitioner, in this case the knowledge engineer, are of direct importance. Although methods such as KADS provide guidance on what types of information are required to instantiate a developing knowledge base, this guidance may not be enough to acquire the cognitively deep reasoning behind expert behaviour which is so necessary to complex decision aids such as the envisaged operational ASuW and ASW systems.

One technique in particular has been used successfully to produce or improve non-KBS decision aids. Gary Klein has developed a model of time-pressured expert decision making, called Recognition Primed Decision Theory (RPD) [Klein and Calderwood, 1991]. It stresses the roles of areas such as situation assessment, mental simulation of earlier events and future event expectancies in expert decisions. It emphasises the recognition component of expert decision making in which experts recognise situations as being similar to past situations and remember the associated solution. While the whole RPD model is described by a flowchart-like representation of activities, much like a KADS Directive Model, it hides a complexity within these activities (such as 'Mental Simulation') which makes it a difficult model to use as a basis for a KBS implementation. One aspect of the model, that of the recognition component, has been used as the basis of Case-Based Reasoning KBS, since the act of recognition and associating a solution is similar to that of matching previously stored cases to situations, associating particular solutions or types of aid to the operator. Such a system has been developed by Noble [1989].

However, Cased-Based KBS techniques are limited in scope and may not be directly applicable to all the activities required of a KBS decision aid. Instead, the model's use of expectations and assumptions, together with an emphasis towards expert situation assessment make it applicable to the generic real-time KBS structure described earlier. Klein has developed an interview technique based around the RPD view of decision making, called Critical Decision Method (CDM). It is proposed that this, or other related critical incident-based interview methods could be used to increase understanding of expert situation assessment and problem solving beyond that achievable through the use of KADS.

CDM uses critical incidents in an expert's experience to understand expert decision making. The technique facilitates the acquisition of the cognitive requirements of a task or skill, rather than just the expert's resultant behaviour. CDM takes an expert back through a critical incident, outlining decision points, information available, expectations and assumptions. By asking particular probe questions within this detailed description of the incident, the expert finds it easier to express the implicit

information behind their decisions. The hypothesis is that the results can be generalised to explain expert behaviour in more common, perhaps less critical scenarios.

It is the ability of CDM to bring out this implicit information that makes it applicable as a technique for the further development of KBS decision aids. Unfortunately, neither RPD or CDM are directly aimed at KA for KBS, but it should be possible to adapt CDM for the generic real-time KBS structure, to elicit the knowledge required for the Situation Assessment module, the expectations and assumptions represented within the structure, and perhaps more importantly to determine how the resulting KBS can best provide aid to the operator.

## 5   Conclusion

The research reported to date has gone a long way to redressing the imbalance between the comprehensive means of representing expertise by the use of the DMUSE real-time distributed software toolkit and the limited capability of the means of capturing the requisite knowledge and expertise. The ASuW KBS Advisor project has proven the concept of using a knowledge acquisition toolkit to design and partially implement (via a translator) a real-time KBS, reducing the knowledge acquisition bottleneck of KBS development. The KADS methodology has been used successfully to structure both the knowledge acquisition activity and the elicited knowledge, using the Directive Models paradigm and the four layer knowledge structure. Whilst this technique is capable of producing initial prototypes, decision aid applications of the complexity envisaged for operational ASuW and ASW systems necessitate the additional use of analytical techniques to more completely understand and encapsulate an operator's situation assessment and decision-making behaviour.
Use of the KADS methodology and automated tools needs to be supplemented by interview techniques such as CDM, which should be adaptable to the KBS development process.

A generalised form of the design of the ASuW KBS Advisor is proposed as a generic structure for real-time KBS decision aids. It provides aid to the operator from both Planning and Situation Assessment activities. The current ASW Advisor (TACAID) is being redesigned using the generic structure. Expansion of its situation assessment has progressed using analytical techniques, producing results which were not apparent from the earlier use of conventional expert interviews.

Importance is placed on developing an accurate representation of operator Situation Assessment, to allow a closer integration between decision aid and operator, and to supply well structured information to KBS planning activities.

## Acknowledgements

## References

[Ellis et al., 1994] R. Ellis, R. Hepworth and H. Howells. Structured and Analytical Knowledge Acquisition Methods For Tactical KBS Decision Aids. Symposium on Command and Control Research and Decision Aids. Naval Postgraduate School, Monterey, California.

[Fransella and Bannister, 1977] F. Fransella and D. Bannister. A Manual for Repertory Grid Technique. Academic Press.

[Howells and Bickerton, 1994] H. Howells and Lt. R. E. Bickerton RN. Decision Making in ASW Helicopters using KBS Advisors. In Proceedings of the First DRA/MOD Workshop in C³I Systems Research. Oxford, UK. (to be published).

[Kelly, 1955] G. A. Kelly. The Psychology of Personal Constructs. Vols 1 and 2. Norton, New York.

[Klein and Calderwood, 1991] G. A. Klein and R. Calderwood. Decision Models: Some Lessons From the Field. IEEE Transactions on Systems, Man, and Cybernetics. Vol 21, No. 5, 1018-1026.

[Noble, 1989] D. F. Noble. Schema-Based knowledge elicitiation for planning and situation assessment aids. IEEE Transactions on Systems, Man, and Cybernetics. Vol 19, No 3, 473-482.

[Wielinga et al., 1992] B. J. Wielinga, A. T. Schreiber and J. A. Breuker. KADS:   A modelling approach to knowledge engineering. Knowledge Acquisition, Special issue on The KADS approach to knowledge engineering. 4(1):5-53.

# DESIGNING REAL TIME DECISION SUPPORT
# FOR FUTURE SYSTEMS AND SCENARIOS

Jeremy Clare, Craig Peden and Ron Smith
Cambridge Consultants Ltd (UK)
Science Park, Milton Road, Cambridge, England, CB4 4DW

## 1 SUMMARY
Future decision support systems will require more intelligent processing capabilities to increase their effectiveness. Knowledge Based Systems offer an attractive solution to meet these decision support roles however, two fundamental problems exist in the knowledge acquisition process which is a key element in their design. The first concerns making best use of contact time with experts, the second relates to the lack of direct expertise in the operation of future, non-mature systems.

REKAP, a knowledge acquisition methodology, offers a solution to these problems through the use of structured analysis and pre-defined models of problem solving inferences. The methodology is discussed in this paper, in particular its value in overcoming the lack of direct knowledge of future system capabilities.

## 2 INTRODUCTION
Real time decision support systems (RTDS) of the future will be characterised by the following trends which highlight the need for more intelligent processing:

- increasing quantities of data will be provided by sensors and other sources.
- this data will require more involved and detailed interpretation.
- there will be reduced time available in which to respond to incoming data.

Such RTDS will be required to deliver greater throughput, quality and consistency than existing systems. This performance improvement from RTDS will be accompanied by improved or new sensors to meet less well defined operational scenarios.

In order to prevent an increase in operator workload, indeed to reduce workload to a more acceptable level, and to achieve the required levels of system performance increasingly sophisticated Decision Support will be required.

In developing an RTDS to support the human crew a key element will be to ensure that the human and machine elements of the overall system act in harmony. Thus it is critical that the decision making process are complimentary. In order to achieve the appropriate harmony we must use development tools which allow an understanding of the relevant roles of the human and the machine.

One of the most promising approaches for providing decision support is the use of Knowledge Based Systems (KBS). KBS technology provides scope for supporting the uncertainty of data and incomplete knowledge of the situation. In addition, KBS applications can be constructed to provide explanation and context which help the individual to reach decisions. However, when we consider the application of KBS to future RTDS, two major areas have to be addressed. Firstly, there is the issue of real time. For a RTDS to function effectively, decisions have to be made in a timely manner. To achieve this, processing must be capable of producing the results when required. In addition, it has to be possible to change the priorities of processing as circumstances change. The second area that has to be addressed is that of acquiring the relevant knowledge needed to build the system. Historically, the Knowledge Acquisition (KA) process has been a bottleneck in development due to the requirement for extensive expert contact time.

In this paper the specific issues of knowledge acquisition for Real Time Decision Support are addressed. Associated work (Martin et al, 1993) has been carried out to enable the development of real time KBS applications which has lead to the development of the experimental D-Muse tool kit. In addressing the KA problem two principal issues have needed to be addressed. Firstly, how to make best use of expert availability and, secondly, how to deal with the limitation in experts' knowledge of how future systems will be used.

In order to address these two fundamental issues of KA a project has been carried out by Cambridge Consultants Limited and the University of Nottingham on behalf of DRA Farnborough. This project has developed the REKAP methodology (Cupit et al, 1993) which seeks to build on the best practice in KA methods and software development methods. To test this methodology, a demonstration application is being developed. In the following sections we discuss the special issues of KA for CIS, outline the REKAP methodology and the tools developed to support it, the discussion of the methodology is illustrated by consideration of a demonstration application.

## 3 KNOWLEDGE ACQUISITION FOR FUTURE CIS APPLICATIONS
The process of knowledge acquisition has long been regarded as one of the more difficult aspects in the development of operational KBS applications. A key part of the problem has been the significant amount of time needed to be spent with domain experts. The number of sessions with the domain expert becomes extensive when there is a need to ensure that the application is complete. This is one of the major differences between the development of illustrative or prototype KBS's and the full blown operational application. To address the problem of extensive contact time a number of major research programmes have been carried out. A key development has been the KADS methodology (Wielinga et al, 1992) which has developed an understanding of how expert knowledge is structured and utilised. This has led to the development of tools to aid the KA process. These are discussed in some detail below.

In the current project the use of a highly structured approach to KA has highlighted the second issue of building future systems. When a future application is being considered then the domain expert faces an additional problem. This is that in some respects he is not a domain expert in the operational function of the system of the future. This means that the domain expert has to reason about how he would utilise the system to solve problems within the

context of the new system operations. This situation inevitably leads to a problem, that different aspects of expertise are elicited with respect to different parts of the problem domain without any necessary internal consistency between the parts.

The reality is that for any future system various sources of expertise needs to be integrated into a coherent whole. The principal sources of expertise are:

- domain experts' knowledge of how tasks and task elements will be performed.
- development experts - knowledge of how system components will perform and how they will interact.
- Operational Requirements experts - knowledge of the roles and operational scenarios in which the system will be deployed.

In the following section we discuss the REKAP methodology and how the structured approach that is adopted helps to integrate the various sources of knowledge.

# 4 BACKGROUND

The REKAP methodology was developed to facilitate the building of KBS for real-time applications. The methodology builds on two earlier developments in the areas of knowledge engineering (the KADS acquisition methodology) and software engineering (real-time structured analysis (SA/RT)). When employed in conjunction these two methodologies complement each other.

## 4.1 The KADS Methodology

Within KADS the central concern of acquisition is the construction of a conceptual model, a description of knowledge possessed by domain experts as opposed to the technical design model, which is a description of the desired KBS.

REKAP is based around extensions to KADS methodology and work arising from the ACKNOWLEDGE (Anjewierden *et al*, 1992) and VITAL (Shadbolt *et al*, 1993) projects.

KADS proposes a four layer structure to "the knowledge level" (Newell, 1982), which is free from any implementational concerns:

- The domain layer. Knowledge within this layer is specific to the problem domain, such as a knowledge of teleologically relevant concepts and relations between such concepts.
- The inference layer. This contains knowledge of the basic inferences involved in problem solving. This is seen as using "meta-classes" of domain knowledge which are independent of the specific domain.
- The task layer. Contains knowledge of the relationship between inferences and tasks. The task layer may be thought of as providing default control flow for inferences.
- The strategic layer. Here meta-control knowledge provides information of how problem solving should be sequenced according to different environmental circumstances.

KADS promotes the idea of 'interpretation models' (Breuker *et al*, 1987)) during the acquisition process. These are abstract, pre-defined models of the inferences of certain types of knowledge, such as situation assessment or diagnosis. Such models also act as abstract templates for expertise (O'Hara, 1993), providing guidance on the types of knowledge required.

A key problem associated with the use of interpretation models is that of initial model selection. More than one model may be applicable to a given function (O'Hara & Shadbolt, 1993a) and,

after re-description, generic functions may become equivalent (O'Hara & Shadbolt, 1993b). To overcome this problem, the ACKNOWLEDGE and VITAL projects introduced the idea of generalised directive models (GDM) (Terpstra *et al*, 1993; Van Heijst *et al*, 1992). These GDM's are progressively refined during acquisition to meet the specific domain needs.

## 4.2 Real Time Structured Analysis (SA/RT)

The CONSENSUS methodology (Bokma *et al*, 1993) was developed to assist in identifying the requirements for KBS applications. CONSENSUS adopts the view of a system specification comprising a requirements model, containing the functionality of the system, and an architecture model, describing how the system will be structured to achieve the functionality.

Within CONSENSUS requirements models are constructed by employing SA/RT techniques, CASE tools can be used to support this process. These provide guidance in the construction of large conventional software systems (Hatley & Pirbhai, 1988) and allow the user to produce hierarchies of data diagrams and process specifications, supplemented by control flow diagrams and specifications.

## 4.3 Demonstration Application

During the REKAP research programme an application was chosen around which to develop and trial the methodology. This application involves providing Decision Support to the crew of a maritime helicopter engaged in Anti-Surface unit Warfare (ASuW). The target aircraft is not expected to have such a DSS facility in the near future, the scenario has thus been set beyond the year 2000.

The application was chosen since it met the real-time and level of complexity criteria needed to test the methodology.

The application scenario, involving the maritime helicopter in the co-ordination of a number of co-operating aircraft, requires that surface contacts are found and monitored while minimising the threat experienced by friendly forces. The surface picture will be complicated by the presence of merchant and other shipping in addition to the hostile forces expected.

The scenario is complex and multi-threaded, involving the assessment and control of a variety of different and dispersed data sources, it also requires reactive actions from the helicopter in order to adjust to the rapidly changing situation.

Knowledge acquisition for the application is complicated, in common with the design of DSS for other future systems, since the mission system in question is not yet in existence. The potential of the platform can only be based on the system specification. In addition, the sensors and systems to be supported will have evolved significantly by the time the DSS may be installed.

# 5 THE REKAP APPROACH

REKAP involves the design of two models; the conceptual model which contains the functional decomposition of the system and is independent of any implementation concerns. The second model, the architectural model, is derived from the conceptual description to meet any specific implementation issues associated with the target software toolkit.

## 5.1 The Conceptual Model

In the identification of the conceptual model a functional description of the CIS is used. This functional description will provide details of the key processes within the system; for example:

- situation assessment,
- data fusion,
- planning,
- mission monitoring.

The identification of a functional description of the system may require the expertise of operational requirements staff rather than system operators. The key is to produce a description of functional roles rather than problem solving and domain knowledge. It is beneficial at this stage to identify which functional areas are best served by the addition of decision support, system operators expertise will be necessary in determining this.

The process of functional description can be illustrated by considering the top level data flow diagram (DFD) produced for the ASuW application in figure 1.

This functional representation now provides the context in which to identify the data needs associated with meeting the various functional goals. With the situation assessment process it will be necessary to have details of the 'fused-picture' and the 'current plan' etc.

The use of this form of structured analysis, supported by CASE tools (in this instance TurboCASE), offers many benefits as the KA activities evolve. Use of TurboCASE offers the means of visualising the complete system, relating the various components and provides consistency checking once more detailed decomposition is achieved.

Having identified the functional description of the system KADS GDM's can be identified which can represent the inference- and task-layer knowledge within the various functions. GDM selection will be influenced by initial KA with experts which will characterise the nature of the problem solving used in the domain. As an example, the GDM in figure 2 was used to model the situation assessment function within the target application.

The GDM provides a means of describing the task structure at subsequent levels in the DFD. The model can be seen to offer a means for representing problem solving knowledge.

Domain experts suggested that situation assessment is based upon the notion of matching the known features of a particular contact (track) with domain specific models of actual objects (schemata). For example, a contact with a speed in excess of 30 knots would be thought unlikely to be a fishing boat or, under most wind conditions, a yacht. However, it may be a warship or a modern merchant ship.

Such reasoning is represented in the GDM where track observables are compared with object definitions to identify matches. The identification of such knowledge is achieved without consideration of the capabilities of the actual system, it is a generic model for solving such problems. The situation assessment GDM is represented within the TurboCASE structure as shown in figures 3 and 4.

Here tracks are selected to be classified. Three major parameters are required; behaviour, identity and group membership. The GDM is used to model the three individual processes. A track is compared with the pre-defined object models and matches recorded. Due to the nature of the domain, matches are never considered absolute, the model tries to refute each possibility rather than confirm. This avoids a contact being classified as friendly where there is even a small possibility of it being hostile.

Expert knowledge is of primary importance in creating the attribute structures of the various objects within the domain. This knowledge evolves in stages, eliciting what is currently possible and eventually accounting for the future potential of the CIS. The final stages involves experts' in some conjecture.

The KA process is focused upon identifying the data requirements of the various processes at all levels in the DFD representation. Within the REKAP method data flows represent objects within the architectural model. These must be structured during the KA activity to provide the object definitions in the final system. Various methods are provided within REKAP to aid this process; laddering, carding sorting, repertory grid construction and analysis, and a rule induction algorithm (CNN). These are available as part of ProtoKEW (Shadbolt, 1992), an automated KA toolset which also offers a means to aid the initial construction of interpretational models.

Within the situation assessment process knowledge was elicitated which allowed various objects to be discriminated, based on different sensor data. This involved experts in providing assessments of future sensor performance however, due to the nature of the REKAP processes such conjecture is constrained. This provides a useful means of considering future system enhancements within the framework of existing capabilities.

The track object produced for the ASuW application is shown in figure 5. The following areas would provide the means for scoping the structure of the object.

- what are the basic attributes of a track?
- how do experts differentiate between tracks?
  - on individual sensors?
  - with integrated data?
- how might they be able to differentiate given a particular new capability or sensor?

Having elicited these various knowledge areas KADS allows the construction of a complete knowledge model. The four layers provide the means of integrating the knowledge:

- data specifications are allocated to the domain layer
- data flows to the task layer
- process specifications to the inference layer
- control to the strategy layer
- control specifications divided between domain, inference and strategy layers.

The relationship between the various knowledge types elicited, when re-created in a target language, provides a complete task and object structure with inference rules and control data as the core of an executable KBS.

## 5.2 Tool Support

The REKAP methodology is supported by a number of software tools. TurboCASE is used to build the functional description of the system. This aids the construction of the task structure

together with creation of the data and control flows within the system. ProtoKEW has already been discussed. This toolset is used to construct the detailed conceptual model.

The use of these tools to support the KA activity provides a means of describing and representing expert knowledge. This offers a powerful aid to visualise the knowledge as it is elicited, both for the knowledge engineer and the expert. The representations within ProtoKEW and TurboCASE provide a common language for discussing the KA requirements.

Providing this additional visual aid within the KA process allows greater participation of experts. This aids the exploration of future system capabilities since consistency of information can be monitored during the KA session.

The tools and the use of the GDM's also provides more focused KA sessions, exploring specific areas of the knowledge model seen as weak or incomplete. The tools also provide traceability of knowledge updates during the evolution of the model.

The final key feature of the REKAP methodology, not discussed in detail in this paper, is the translation of the conceptual model into run-time implementation code. REKAP achieves this via a translator tool which automatically takes the conceptual model and maps it, using the KADS four layer model, into the Muse real-time AI software toolkit.

## 6 CONCLUDING REMARKS

This paper has attempted to demonstrate how REKAP, through the use of Generalised Directive Models and structured analysis, offers a powerful methodology to overcome the difficulties associated with designing KBS decision support tools. In adopting a functional approach to the design of the KBS architecture the refinement, during the knowledge acquisition phase of design, of the GDMs provides a means for incorporating the supposition of experts regarding future system potential.
In support of the REKAP approach it is necessary to:
- Provide a functional description of the system, either directly from specifications and requirements studies, or from knowledge acquisition with Operational Requirements staff.
- Identify problem solving models applicable to the key functions to be supported by the DSS through KA with operators.
- Employ principled knowledge acquisition methods to identify the structure and content of knowledge at each of the KADS four levels.
- Gradually refine the knowledge acquired to include theoretical information regarding future sensor and system enhanced capabilities.

REKAP, and the tools which support it, provide a framework in which to visualise and maintain the acquired knowledge during the initial KA phase of design, providing consistency checking and other housekeeping functions. The tool support has also been seen to improve the productivity of KA sessions by providing the expert and knowledge engineer a means by which both can view and understand the data acquired.

We believe these benefits, together with the automatic translation of captured knowledge to the Muse AI toolkit make the REKAP approach a powerful and valuable method of designing decision support tools for future CIS.

## References

Anjewierden, A., Wielinga, B. and Shadbolt, N. (1992) Supporting Knowledge Acquisition: The ACKnowledge Project. In Steels, L. and Lepape, B. (Eds.) ESPRIT-92 Knowledge Engineering, CEC, Brussels.

Bokma, A., Slade, A., Bateman, M., Kellaway, M. and Martin, S. (1993) The CONSENSUS Method, Final Report D17. IEATP Project 1365, British Aerospace.

Breuker, J., Wielinga, B., Van Someren, M., De Hoog, R., Schreiber, G., De Greef, P., Bredeweg, B., Wielemaker, L., Billault, J-P., Davoodi, M. and Hayward, S. (1987) Model Driven Knowledge Acquisition: interpretation models. ESPRIT Project P1098 Deliverable D1. University of Amsterdam and STL Ltd.

Cupit, J., Major, N., Shadbolt, N., Smith, R. and Clare, J. (1993) REKAP: A Methodology for the Automated Construction of Real-Time and Distributed Knowledge-Based Systems. To be published at 14th International Avignon Conference, Paris (1994).

Hatley, D. and Pirbhai, I. (1987) Strategies for Real-Time System Specification. Dorset House Publishing, New York.

Martin, S., Zanconato, R. and Smith, R. (1993) Architectures for Practical Real-Time Knowledge Based Systems. In Proceedings of Expert Systems (1993), British Computing Society, pp 83-97.

Newell, A. (1982) The Knowledge Level. *Artificial Intelligence*, 18, pp 87-127.

O'Hara, K. (1993) A Representation of KADS-I Interpretation Models Using a Decompositional Approach. In Lockenhoff, C., Fensel, D. and Studer, R. (Eds.) *Proceedings of the 3rd KADS Meeting*, pp 147-169. Siemens AG, Munich.

O'Hara, K. and Shadbolt, N. (1993a) AI Models as a Variety of Psychological Explanation. In the *Proceedings of IJCAI-93*, pp 188-193.

O'Hara, K. and Shadbolt, N. (1993b) Locating Generic Tasks. *Knowledge Acquisition*, 5(4).

Shadbolt, N. (1992) Facts, Fantasies and Frameworks: the design of a knowledge acquisition workbench. In Schmalhofer, F., Strube, G. and Wetter, T. (Eds.) Contemporary Knowledge Engineering and Cognition. Springer Verlag, Heidelberg.

Shadbolt, N., Motta, E. and Rouge, A. (1993) Constructing Knowledge-Based Systems. *IEEE Software, November 1993*, pp 34-39.

Terpstra, P., Van Heijst, G., Shadbolt, N. and Wielinga, B. (1993) Knowledge Acquisition Process Support Through Generalised Directive Models. In David, J-M., Krivine, J-P. and Simmons, R. (Eds.) *Second Generation Expert Systems*, pp 428-454. Sringer-Verlag.

Van Heijst, G., Terpstra, P., Wielinga, B. and Shadbolt, N. (1992) Using Generalised Directive Models in Knowledge Acquisition. In Wetter, T., Althoff, K-D., Boose, J., Gaines, B., Linster, M. and Schmalhofer, F. (Eds.) *Current Developments in Knowledge Acquisition - EKAW 92*, pp 112-132. Springer-Verlag.

Wielinga, B., Schreiber, A. and Breuker, J. (1992) KADS: A Modelling Approach to Knowledge Engineering. *Knowledge Acquisition*, 4(1). Special issue 'The KADS Approach to Knowledge Engineering'.

**Figure 1 - 1st level data flow decomposition of the ASuW DSS**



**Figure 2 - The GDM for situation assessment**

**Figure 3 - Track analysis data flow decomposition**



**Figure 4 - Behavioural analysis data flow decomposition**



**Figure 5 - Track object structure**

# SESSION III - TRUST DEVELOPMENT

## SYNOPSIS

This section comprises papers on the trustworthiness of HE-C decisions. Papers 11 and 12 discuss trust from experience with technology demonstration systems; Papers 13 and 14 report the results of psychological investigations into human-human and HE-C trust, with reference to a teamwork model; Paper 15 reviews the role of memory in trust development; and Paper 16 discusses EC intelligence, and proposes guidelines for building trust. *Paper 11* is based on trials with a KBS system for naval command and control. Measures of trust point to the importance of the system's accuracy and predictability. To be trusted, the system must demonstrate competent role performance, and provide facilities that enable operators to predict its accuracy. To accept a DSS hypothesis, operators need to understand and agree with the DSS reasoning process, with easily comprehended explanation. *Paper 12* discusses an AI warning system for civil aircraft. In emergencies, the need to trust information is raised by time pressure, and by the complexity of simultaneous multiple failures. Trust is achieved by basing the design on user models of the system operation using Model-Based Reasoning (MBR), and by producing information that reflects the goals and expectations of the user. Primary causes of failure are isolated using understandable fuzzy reasoning logic. Other trust enhancing features include: a goal hierarchy for the generation of logical prioritised actions; a minimal set of ordered actions consistent with checklist procedures; and anticipation of action consequences to indicate that the system is responding intelligently. *Paper 13* reports a study of simulated aircraft adaptive automation (i.e. EC) using the Multi-Attribute Task (MAT) battery. Failures were introduced in the ability of the EC to provide levels of aiding in a timely and appropriate manner. Subjects compensated for poor EC performance without awareness of adaptation failure. Trust ratings were related to perceptions of EC competence, but not adaptation failure. Performance was associated with rated situation awareness. Procedural safeguards, such as prime directive protocols, are needed against the consequences of inappropriate task allocations and undetected adaptation failures, and to prevent false trust being engendered in imperfect adaptive aiding. Audit data of teamwork quality are presented based on a model of teamwork goals, resources, structure and processes. *Paper 14* reports the results of a communications analysis of recordings from Tornado aircraft simulator missions. Communications were categorised as statements, assertions, questions, confirmations or discussions. More unprompted statements were made than any other form of vocal communication, particularly by the navigator, to maintain common knowledge and situation awareness. The data show a two-way flow of information and initiative taking. The absence of contradictions and the small number of questions are indicative of a high degree of trust. Conclusions are drawn for HE-C communication with reference to the above teamwork model. For EC to take over the navigator's functions, this pattern of bi-lateral communication will need to be reproduced, with intent inferencing and common knowledge to support shared initiative taking. *Paper 15* considers how communication, considered as the sharing of knowledge, involves remembering. It reports psychological data on memory enhancing techniques, in particular the Cognitive Interview (CI). CI is a proven communication aid in a range of contexts, including market research and trauma victim interviews. The authors consider how features of communication, such as trust and rapport building, can be improved, say during HE-C mission debriefing, by the application of memory techniques such as CI. *Paper 16* considers if the HE-C team would engender more trust if both team members appeared human. The Turing test is described which determines whether an interrogator can tell the difference between human and machine answers to questions. If not, the machine possesses the qualities of intelligence, and can be trusted. Weaknesses of the test are identified. Guidelines are proposed for building trust, based on the need for consistency and correctness in team decisions. These include prime directives, levels of autonomy, conformance with the pilot's mental model, transparent interfaces, and summarised information. A manager-staff relationship is proposed as the ideal team structure. Internal trust will lead to efficient, consistently correct performance, which will engender trust in the team from others.

# Seaworthy Trust: Confidence in Automated Data Fusion

A.Simpson
G.N.Brander

DRA Portsdown
Portsmouth
Hampshire
PO6 4AA
United Kingdom

## 1. SUMMARY

If the Human-Electronic Crew is to function effectively as a trustworthy team, an issue of great importance is the degree to which the Human component of the Team trusts the Electronic one. In the context of a Knowledge Based System (KBS), the operator may constantly check the system's output, or he may always accept it without any basis for so doing. In either case, the Team is not functioning effectively.

Trust in a KBS was examined in a study of users of the Data Fusion Technology Demonstrator System (DFTDS). This is a prototype Command and Control ($C^2$) system intended to explore the capabilities of Knowledge Based System techniques in picture compilation at sea through the automated presentation of fused information.

The results of the study indicated that the system's accuracy and predictability were important. Predictability is influenced by operators' comprehension of the DFTDS, which is affected by the relationship between the decision-making process employed by the system and the operator, its ability to provide explanations of this process, the content and format of the explanations, and the way in which the DFTDS deals with the uncertainty inherent in its hypotheses. In order that a system can be trusted, it must both demonstrate technically competent role performance and provide its operators with facilities that enable them to predict the pattern of its accuracy.

## 2. INTRODUCTION

In an interaction with a KBS, there are two possible extremes of response by the operator: he always accepts the system's hypotheses with or without any basis for doing so (blind trust), or he questions all of them, with or without any basis for doing so. The former is obviously undesirable, as we do not have infallible systems, and so the operator needs to be kept 'in the loop' to cope with those situations which the system is unable to cope with (perhaps because they have not been foreseen). Questioning all of the system's hypotheses is equally undesirable, and may actually increase the operator's workload or degrade his performance in comparison to that attained when not using the KBS. Overconfidence may be less desirable than underconfidence, as the consequences are potentially more harmful.

It is therefore important that operators should have an appropriate level of trust in a KBS (Ref 1). This might be interpreted as a willingness to accept its hypothesis without question in situations/for functions which it is known to be capable of handling at least as well as (if not better than) the operator himself, and questioning the system in situations/for functions which it cannot handle as well as the operator.

Factors which may influence the operator's trust in a KBS include the perceived accuracy of the system, its predictability, the provision of explanation facilities, the decision-making model used by the system and the way in which it deals with uncertainty. The provision of explanation facilities, the decision-making model used by the system and the way in which it deals with uncertainty affect the operator's trust in the system because they influence the degree to which he is able to understand and evaluate the reasons for the system's actions and determine whether or not it presents a true picture of the situation.

### 2.1 Trust in the DFTDS

The role of automated support within a $C^2$ system is to provide the right information, at the right time, to the right user, and in a form that can be assimilated and acted upon, especially in time-critical situations. In present Operations Room procedures, the operator has to trust those above and below him for providing information in the chain of Command. A critical issue, therefore, is an operator's trust in the outcome of the DFTDS's data fusion process.

### 2.2 Explanation Facilities

It is generally agreed that explanation facilities are a distinguishing feature of KBSs (Ref 2). They are perhaps especially important during early usage of the system, when the operator may build up trust in the system and determine its limits. If the operator finds that the system's conclusions are well-supported and comparable with, or better than, his own mental processing, then he is more likely to come to believe the system without asking for justification. Explanation facilities may also allow the operator to query the system during periods of low activity or training exercises, so generating trust in it for time-critical situations in which there is no opportunity for querying the system's hypothesis.

Buchanan and Shortliffe (Ref 2) consider that explanation facilities are crucial to the success of a KBS, and perform the following functions: assisting users and system builders in understanding the contents of the system's knowledge base and reasoning processes; facilitating the debugging of the system during development; educating users about the domain and capabilities of the system; and persuading users that the system's conclusions are correct, so that they can ultimately accept these conclusions and trust the system's reasoning powers.

According to Abu-Hakima and Oppacher (Ref 3), genuine explanations show why an action is reasonable in the light of available information, heuristics and domain principles, and why an alternative action was not performed. The authors contrast genuine explanations with the type of explanation in which the system merely provides a trace of

the rules used in order to arrive at a hypothesis. As Kidd (Ref 4) points out, the latter may be sufficient for system debugging by the knowledge engineer, but it does not necessarily provide the operator with the kind of information that he is seeking in order to judge whether or not to accept the system's hypothesis.

Although there is some evidence that explanation facilities are under-used and perhaps not even necessary in some systems (Ref 5), this may be because the explanations provided by the system do not meet the user's needs. Alternatively, explanations may only be required for certain tasks. At present, little is known about what constitutes an acceptable explanation for users of real-time KBSs.

The extent to which operators access the explanation facilities provided by the DFTDS may be an indication of the degree to which they trust the system, i.e., do they accept its hypotheses without question? The type of information being sought by operators may include: why/why-not/how evidence was used; why/why-not a conclusion was reached; how a solution was arrived at. However, this is a complex issue, in that the extent to which the explanation facilities are used may be influenced by their content, the way in which this is presented to the operator, and the ease with which it can be accessed.

In the software vesion used during the two trials periods, there were eight types of explanation provided by the DFTDS, on separate "pages" arranged in a hierarchy within the explanation window. In the light of experience the explanation facility is being redesigned for the next software release.

When the DFTDS was initially designed, it was envisaged that the explanations would be mainly accessed by the Command. They were also intended to provide the system designers with feedback concerning the performance of the data fusion module within the DFTDS. The present study attempted to determine the extent to which the explanations were accessed by operators, the type of information that operators were seeking, and whether all of the information that they required was available.

## 2.3 Decision Making Model

Another important factor is the decision-making model employed by a KBS, as this influences the user's understanding of the system, and hence his trust in it.

The operator will only be confident in the system's hypotheses if the reasoning processes it employs are readily intelligible to him. This does not necessarily mean that the system has to possess a psychological model which exactly imitates the human's reasoning process, but it does mean that the representation must be able to capture the range and power of the human expert's knowledge in the particular domain.

Thus if the expert uses or recognises key features in the information environment, the system should also make use of these in the same way and in the same order.

If a KBS does 'reason' in the same way as the operator, then it is more readily able to produce an explanation which the user can understand. In addition, it is more likely that which the information available from the system is that which the operator would use to make a decision for himself, and for it to be presented in the required order. This means that it is easier and quicker for him to verify a decision if required. The counterargument is that a KBS is intended to enhance overall system performance, but in order to accomplish this aim, it may well be that tasks have to be performed in a

different way to that used by operators at present. It is not as yet known whether it is necessary for the system to perform a task in the same way as the operator would, or if it is only the outcome which is important.

Most of the research to date has focussed on medical diagnosis systems, but there is a critical difference between these and Naval $C^2$ Systems, in that the pace is real-time in the latter.

The present study attempted to determine whether there were any differences between the reasoning processes used by operators and the DFTDS.

## 2.4 Representation of Uncertainty

The hypotheses produced by the DFTDS have varying degrees of uncertainty associated with them. This is due to attributes of the sensor data used by the DFTDS, and the fusion process itself. The system only displays its best hypothesis, but others are constructed in the machine.

The way in which the DFTDS represents and conveys the uncertainty inherent in its hypotheses may influence operators' trust in the system, in that it affects their comprehension of its reasoning process.

A critical issue is whether the operator should be, or needs to be, aware of this uncertainty. It has been proposed (Ref 6) that an awareness that uncertainty exists is crucial for an accurate understanding of the situation, and there is some evidence that the provision of probability information increases the operator's confidence in a system (Ref 7). If this is the case, then there is a need to determine how this uncertainty should be represented to the operator.

Uncertainty is represented in the DFTDS in two ways (i) the degree of certainty with which the system determines the Standard Identity of a platform is displayed numerically by the allocation of penalty points, which can be seen in one of the explanation windows, and (ii) limited information is shown on the tactical display in the form of a single letter in the Track Label. The latter provides information such as when there is an ambiguous correlation associated with a vehicle.

The present study attempted to determine whether operators required to know the degree of certainty with which the DFTDS held a hypothesis, and if this was represented in a readily comprehensible format.

## 3. METHOD

The DFTDS was not designed to support specific job-related tasks – the intention was to see what operators used the available technology for. Consequently, the study focussed on the ways in which the DFTDS was used, rather than examining it during a set of predetermined user tasks.

There were two main trials periods: June 1992 (Trial 1) and February 1993 (Trial 2). The most important set of data was that obtained in Trial 2. This was the longest exercise in which the DFTDS was used, and operators were more experienced in using the DFTDS than they had been during previous exercises.

Data was collected by means of questionnaires, semi-structured interviews, video recordings, and software logs.

## 4. RESULTS AND DISCUSSION

In Section 2, it was proposed that operators' trust in a KBS may be influenced by its perceived accuracy, the predictability of the system and factors which affect operators' comprehension of the system. The latter include the provision and content of explanation facilities, the

relationship between the decision-making model it employs and that used by the operator, and the way in which the system deals with uncertainty.

## 4.1 Trust
Operators' trust in the DFTDS was measured by the analysis of subjective data obtained from questionnaires and interviews and objective data from software logs.

### 4.1.1 Subjective measures
Responses to questionnaires completed after Trial 1 indicated that 50% of operators trusted the DFTDS, 25% did not and the other 25% did not know whether they trusted it or not. After Trial 2, 40% of operators trusted the DFTDS, 20% did not trust it, and the remaining 40% were unsure if they trusted the system or not. The general trend was for operators to become more unsure over time as to whether or not they trusted the system.

Sheridan and Hennessey (Ref 8) found that operators, particularly novices, were biased towards distrust in a supervisory control environment, but the present data is not in agreement with this. One possible reason is that operators of the DFTDS, although relatively new to the system, were not novices in their jobs.

During interviews conducted after Trial 1, operators made a number of comments concerning factors which they considered to influence their trust in the DFTDS.

Some operators said that they were taught not to trust a computer, but if they had access to an explanation of its decision-making process, they could see whether they agreed with the process or not. This implies that operators feel it is important to see the way in which the system arrives at a hypothesis, and also to agree with the reasoning process it uses in order to reach this hypothesis.

One operator said that as long as he knew why the DFTDS was wrong, then he trusted it. This again points to the importance of the operator understanding the decision-making process employed by the DFTDS. However, in section 2.3, it was stated that it is not known whether the decision-making process is important, or merely the outcome of this process. The above comment would appear to suggest that it is the process which is the most important factor, not its outcome.

Collins (Ref 9) refers to two types of explanations: felicitous and refutable. The former merely makes the operator feel happy about a system's hypothesis, but the latter places him in a position whereby he is able to agree with the decision-making process or not. The opinion that the system is acceptable if an explanation is available, regardless of whether the DFTDS is correct or not, suggests that it is the ability of the operator to refute a hypothesis that is important, and is therefore linked to the question of what information should be contained in explanations provided by the DFTDS.

The DFTDS was compared to existing systems and equipment, for example, one operator said that he saw no reason to question the DFTDS when he accepted what other Systems told him – he was not able to question them. This may be one reason why operators did not access the explanations available in the DFTDS: they were not used to having the information available. One the other hand, it may be that other systems do not require such a facility because they do not perform any Knowledge Based functions.

Operators' initial expectations were that the system would work, but it was emphasised that if these expectations were not fulfilled, then distrust would quickly occur, thus implying that trust is dependent upon accuracy.

In summary, operators considered the following factors to be important, and to influence their trust in the DFTDS: prior expectations, the decision-making process it employs, the accuracy of the system, its predictability and the provision of explanations. However, none of these factors were found to be related to the subjective measure of trust. It is recommended that this question be further examined in a controlled setting, because factors such as changes in personnel, differences in operator training and poor sensor performance (which affected the data fusion process) also influenced operators' interactions with the DFTDS.

### 4.1.2 Objective measures
It was proposed that the type and frequency of operators' contributions to the data fusion process would indicate their degree of trust in the hypotheses provided by the DFTDS. An analysis of the software log files for Trial 2 showed that manual correlations and decorrelations were indeed made by operators.

The two main reasons given by operators for making manual correlations were (i) because the operator had access to information that the DFTDS did not, and (ii) the operator was forcing correlations in order to clear up the tactical display. Decorrelations were mainly made on the basis of information which was available to operators, but not to the DFTDS.

Operators did not accept the hypothesis proposed by the DFTDS on all occasions. Disagreement with the DFTDS's hypothesis is mainly represented by the changes made to Standard Identity (i.e., hostility) and Platform Identity (i.e., type): 818 attempts were made to change Standard Identity, and 648 to change Platform Identity.

All users for both trials said that they would change a hypothesis made by the DFTDS although, in general, operators said that they would need to be very certain that the DFTDS was wrong before they would make any change.

Reasons given by operators for changing the DFTDS's hypotheses fell into two main categories (i) operators cited specific instances of when they did not agree with the DFTDS's reasoning (e.g., platforms travelling at a speed of 25 knots being labelled as aircraft rather than surface ships), or (ii) more general, ill-defined, reasons (e.g., if the operator "knew" that the DFTDS was wrong).

A number of operators said that they had access to more information than the DFTDS (e.g., Officer of the Watch, GDP, voice, or visual sightings). One operator expressed the opinion that there are almost always factors which are not taken into account by the DFTDS in its decision-making process, because there is no way that they can be represented in the DFTDS.

However, the majority of operators were unable to specify why they would change a DFTDS's hypothesis, beyond saying that it was because of a 'gut feeling' or similar expression. This is considered likely to be due to the reliance of human operators on contextual information.

In summary, operators contributed to the data fusion process because they had access to information that the DFTDS did not, they applied different rules to the data (e.g., different weightings), or their experience lead them to believe that the DFTDS was incorrect. An important point was made by those operators who said that they would make changes to the DFTDS as a result of their own knowledge of

74

the history of the current situation and their experience of how situations had typically developed before and how the tactical picture had built up.

## 4.2 Accuracy

Responses to questionnaires completed after Trial 1 indicated that 50% of operators thought that they knew how accurate the DFTDS was, and the other 50% did not. After Trial 2, 90% of users thought that they knew how accurate the DFTDS was, and only 10% did not. This result was to be expected -- increased usage of the system enabled operators to build up a model of the system's performance in terms of its accuracy. However, subjective knowledge of its accuracy did not appear to significantly increase or decrease operators' trust in the system: the more sure they were that they knew how accurate the DFTDS was, the more unsure operators were whether or not they trusted it. This suggests that factors other than knowledge of a system's accuracy influence operators' trust in the system.

## 4.3 Predictability

After Trial 1, 75% of operators were unsure whether the DFTDS was predictable or not, and the other 25% thought that it was predictable most of the time. After Trial 2, 80% of operators were unsure if the DFTDS was predictable and the other 20% considered that it was not always so.

It would appear that increased usage enabled operators to form a clearer picture of the system's accuracy, but not the pattern of this accuracy – operators claimed that they knew how accurate the DFTDS was, but they were unable to predict the occasions on which it would not be accurate. Increased judgements of unpredictability may reflect a more accurate model of the system – it has been shown that the DFTDS will not always arrive at the same hypothesis given the same input data.

According to Muir (Ref 1), in the early stages of an interaction, trust is based on predictability. The latter is assessed by recurrent behaviours, and trust develops when the operator is able to estimate the system's predictability. The author also speaks of an operator's ability to calibrate his trust in a system, and thus it may not be accuracy *per se* which is the important factor, but the ability of the operator to predict when the DFTDS will, and will not, be correct. The fact that operators were unable to do this means that they do not have calibrated trust in the system.

The concept of predictability does not in itself imply performance accuracy – a system could be predictably incorrect. If operators were equating trust with predictability, then it may be that a system can be thought to be trustworthy even if it produces an incorrect hypothesis. Nonetheless, it is proposed that to trust a KBS means that it must be considered to produce an acceptable output, and so a system that produces an incorrect hypothesis does not fulfil the criterion for trustworthiness. This proposition is in agreement with Muir's (Ref 1) opinion that technically competent role performance is the most influential factor in the development of trust. He further suggests that each person has a 'criterion of competence', below which a system will be judged untrustworthy. This has important implications for the use of Knowledge Based technology in Naval $C^2$ Systems, in that the criterion of competence, and hence trust, may vary between individual operators.

## 4.4 Explanation Facilities

Subjective data concerning the usage of the DFTDS's explanation facilities was obtained from questionnaires and interviews, and objective data by the analysis of software logs and video recordings.

After Trial 1, 75% of operators said that they often asked for an explanation of the DFTDS's decisions, and the other 25% did not do so very often. In the post Trial 2 questionnaires, explanations were said to be used very often by 40% of operators, often by 50%, and not very often by the other 10%.

In section 4.1, it was noted that operators considered that the availability of explanations was an important influence on their trust in the DFTDS. However, an analysis of the software log files and video recordings showed that, in contrast with their subjective opinions, the explanations provided by the DFTDS were rarely accessed by operators. There are several possible reasons for this: operators experience difficulty in accessing the explanations; there is insufficient time to obtain an explanation from the system; the explanations do not contain the information they require; they just "know" that a hypothesis is incorrect, and are not interested in questioning the reasoning used by the system to find out why it is wrong; the information is not presented in a form that is readily comprehended by the operator.

Subjective data does not appear to support the first possible reason. After Trial 1, 25% of operators felt that obtaining the information that they required from the explanations was very easy and 75% felt that it was alright. Responses to post Trial 2 questionnaires indicated that accessing the information contained in the explanations was thought to be easy by 30% of operators, alright by 60%, and difficult by 10%. However, operators did remark that it took a long time to get to information at lower levels of the hierarchy of explanation pages.

The second possible reason why the explanations provided by the DFTDS were rarely accessed is to some extent supported by operators' comments in interviews and questionnaires. Indeed, when the system was designed, it was anticipated that there would not always be sufficient time for operators to access the explanations in a time-critical situation, and that they may be used to build up trust during periods of low activity.

After Trial 1, none of the operators thought that there was any information they would use to make a decision, that was not provided by the DFTDS's explanations. However, after Trial 2, information not contained in the explanation facilities was said to be required by 38% of respondents. Very few operators responded when they were asked what additional information they would use to make a decision. The majority of them said that they would use their own instinct, built up through experience.

The way in which information is presented in the explanations provided by the DFTDS may influence the frequency with which operators access them. If the information is not presented in a form which is easily and quickly comprehended, then operators will be less likely to access the explanations. Operators' comments in interviews conducted in June 1993 provide some support for this proposition, particularly in relation to the way in which the DFTDS represents the uncertainty associated with its hypotheses (see section 4.6).

In summary, the data obtained indicates that the explanations provided by the DFTDS were rarely accessed during Trial 2. It is proposed that this may be because the explanations do not contain the information that operators require, it is not presented in a form that they can readily comprehend, or operators just 'know' that a hypothesis is incorrect and are not interested in the reasons why. It is likely that the present data represents a combination of these reasons, and more research is needed in order to

clarify the issue. A major problem would seem to be that the explanations were primarily intended to assist developers in examining the performance of the data fusion components of the system, rather than for operators use.

## 4.5    Decision Making Model

Subjective data concerning possible differences between the reasoning processes used by operators and the DFTDS was obtained from interviews and questionnaires.

After Trial 1, 75% of operators considered that the DFTDS sometimes "thought" in the same way as they did, and the other 25% considered that it never did. After Trial 2, 70% of operators said that the DFTDS sometimes "thought" in the same way as they did, and the other 30% said that it never did. Increased usage did not, therefore, alter operators' responses to this question, and it would appear that the decision-making process employed by the DFTDS does not accurately reflect that used by the operators for all situations.

After Trial 1, 50% of operators said that they would sometimes accept a hypothesis if they did not agree with the reasoning behind it, and the other 50% never would. After Trial 2, 60% of operators would sometimes accept a hypothesis if they did not agree with the reasoning used by the DFTDS to arrive at that hypothesis, while the other 40% would never do so.

One operator made an important point when he said that the tactical situation and the possible consequences of an incorrect hypothesis would strongly influence his willingness to accept a hypothesis made by the DFTDS if he did not agree with the reasoning behind it. This opinion would seem to support Muir's (Ref 1) proposition that KBSs have an element of risk associated with them, and hence tend to be initially mistrusted. There is not usually much risk associated with human-human relationships, and so the initial tendency is for trust.

When asked to give examples of times when there was a mismatch between they way in which they would reason and the way in which the DFTDS did so, operators either cited specific instances relating to the speed of objects on the display, or gave examples of when they would have a different hypothesis (e.g., 2050 contacts always being shown as "unknown" by the DFTDS, whereas the operator would rather that they were displayed as "suspect" unless proved otherwise).

Differences between the way in which the DFTDS performs data fusion and the way in which operators carry out the process were found by Sherwood-Jones and Northcote (Ref 10). Operators may use different criteria than those used by the DFTDS or give the criteria different weightings, which may vary according to the situation (which they do not in the DFTDS). They also have access to richer information, from a wider range of sources, than the DFTDS does.

In conclusion, the data suggests that the decision-making process employed by the DFTDS is not the same as that used by its operators, and nearly half of those questioned stated that they would never accept a hypothesis proposed by the DFTDS if they did not agree with the system's reasoning process. This issue requires further examination, in order to determine the nature and extent of differences between the decision-making processes used by the DFTDS and its operators, and their impact on operators' trust in the system and acceptance of its output. Sherwood-Jones and Northcote consider that the DFTDS's decision-making should be based on that used by the operator. This proposition must obviously be investigated in future studies.

## 4.6    Representation of Uncertainty

Subjective data was obtained from interviews and questionnaires, and objective data from the analysis of software logs and video recordings.

It was found that the "explanation" window showing the penalty points associated with a hypothesis was not accessed at all during those periods of Trial 2 when operators' interactions with the DFTDS were recorded. In interviews, operators expressed a requirement for a more comprehensible representation of the degree of certainty held by the system, preferably in a graphical format. There is some evidence that although numerical probability levels may provide a sophisticated tool for system designers to reason with uncertainty, they are not meaningful to the users of a KBS (Ref 11).

Operators also considered that it was important to show not only the parameters used and the total "penalty points" for possible identities, but also the numerical contribution made by each parameter to the total number of points. This was seen to be of assistance in an appreciation of their relative weighting factors. One operator remarked that both the inputs to the DFTDS and his own valuations incorporate weighting factors. His are built up through experience, knowledge and prejudice, but he has no visibility of weighting factors used by the DFTDS.

It would therefore appear that there are questions to be answered in relation to the way in which uncertainty is represented in a KBS, how it is explained to the operator, and the relationship between the way in which the system and its operators deal with with that uncertainty. For example, it is not known under what circumstances the operator may be able to cope with multiple hypotheses, or how these and their associated degrees of uncertainty should be represented. Such questions must be addressed in future research.

## 5.    CONCLUSION

The data obtained suggested that users of the DFTDS had neither blind trust in the system nor questioned all of its hypotheses. Its reasoning was sometimes questioned through the use of the explanation facilities provided by the DFTDS (although they were not accessed as much as operators' subjective opinions would suggest), and operators made manual inputs to the data fusion process, thereby changing the hypothesis produced by the system. Manual inputs were made because operators had access to information that the DFTDS did not, they applied different rules to the data, or their experience lead them to believe that the DFTDS was incorrect. There would appear to be differences between the decision-making processes employed by the system and its operators, especially in the way in which they deal with uncertainty.

The results of the study indicated that accuracy and predictability were the most important factors. Predictability is influenced by operators' comprehension of the DFTDS, which, in turn, is affected by the relationship between the decision-making process employed by the system and the operator, its ability to provide explanations of this process, the content and format of the explanations, and the way in which the DFTDS deals with the uncertainty inherent in its hypotheses. In order that a system can be trusted, it must both demonstrate technically competent role performance and provide its operators with facilities that enable them to predict the pattern of its accuracy (cf. Muir (Ref 1)).

It is interesting to note that Muir and Moray (Ref 12) found that a small variable error had the same effect on operators' trust in a system as did a large constant one. The first

affects the system's predictability and the second its accuracy.

## 6. FUTURE RESEARCH

The effectiveness of the Human-Electronic crew is influenced by the degree to which the human component of the team trusts the electronic one.

The data obtained from the present study suggests that if the operator is to trust a KBS, then it must not only demonstrate a level of performance that satisfies the operator's criterion of competence, but also provide the necessary facilities for him to predict those occasions when the system will, and will not be correct.

It is therefore proposed that future research should address the HCI issues that influence operators' ability to predict the performance of a KBS. If they are to be able to predict a system's performance, operators need to understand its decision-making process and, in particular, the way in which it deals with, and represents, uncertainty. The system must be able to explain its reasoning process to the operator, providing him with the information he requires in an easily comprehended format.

Operators also need to understand the reasoning process employed by a KBS in order to make inputs to the data fusion process. These will be required for the forseeable future, as we do not have the technology to build infallible systems. If the operator is to understand the process used by the system, he must be provided with explanations which give him the information that he needs in a format he can readily comprehend and access quickly.

If such explanations are to be provided, the way in which the system represents and reasons with uncertainty is important, and also the relationship between the way in which the operator and the system reason about the information available to them – the decision-making process used in order to arrive at a hypothesis.

The data obtained suggests that the decision-making process employed by the DFTDS is not the same as that used by its operators. Further work is required in order to determine the nature and extent of the differences, and their impact on operators' trust in the system and acceptance of its output. This may be particularly important in view of the fact that many operators said that they would never accept a hypothesis proposed by the DFTDS if they did not agree with its reasoning process. In addition, it is not known how much decision-making processes vary between individual operators.

Further examination of the way in which uncertainty is dealt with and represented in a KBS is also required. Under what circumstances is the operator able to cope with multiple hypotheses, how many alternatives should be shown, and how should these and their associated degrees of uncertainty should be represented to the user?

## 7. REFERENCES

1. MUIR, B.M. Trust between humans and machines, and the design of decision aids. *International Journal of Man-Machine Studies* 27, 527-539. 1987.

2. BUCHANAN, B.G. and SHORTLIFFE, E.H. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Addison-Wesley Publishing Company. 1984.

3. ABU-HAKIMA, S. and OPPACHER, F. Improving explanations in knowledge based systems: RATIONALE. *Knowledge Acquisition* 2(4), 279-390. 1990.

4. KIDD, A.L. The consultative role of an expert system. In: Johnson, P. and Cook, S. (eds.). *People and Computers: Designing the Interface*. Cambridge:Cambridge University Press. 1985.

5. JOHNSON, L. Competent expert systems: A framework for more adequate explanations. *IKBS Workshop on 'explanation'*. University of Surrey. IEE. 1986.

6. SELCON, S.J. and TAYLOR, R.M. Decision support and situational awareness. In: Taylor, R.M. (ed.). Situational Awareness in Dynamic Systems. IAM Report No.708. Proceedings of the 11th Congress of the International Ergonomics Association, Paris, France. 1991.

7. SELCON, S.J., SHADRAKE, R.A. and TAYLOR, R.M. An experimental investigation of explicit probability information as a decision support methodology for future airborne systems. IAM Report No. 692, RAF Institute of Aviation medicine, Farnborough, Hants. 1990.

8. SHERIDAN, T.B and HENNESSEY, R.T (eds.). *Research and Modelling of Supervisory Control Behavior: Report of a Workshop*. Washington, DC: National Academy Press. 1984.

9. COLLINS, H.M. The concept of explanation in expert systems. In: *IKBS Workshop on 'explanation'*. University of Surrey. IEE. 1986.

10. SHERWOOD-JONES, B.M. and NORTHCOTE, E.A. TDS user knowledge base interface validation: Knowledge acquisition with COPE. DRA Report TDP/6.11/1.

11. LAMBERTI, D. and WALLACE, W.A. Presenting uncertainty in expert systems: An issue in information portrayal. *Information and Management* 13, 159-69. 1987.

12. MUIR, B.M. and MORAY, N.P. Operator's trust in relation to system faults. *IEEE Transactions on Systems, Man and Cybernetics*. 1987.

C. R. Ovenden and A. F. Starr

Research and Product Technology
Smiths Industries Aerospace and Defence Systems
Cheltenham, Glos, GL52 4SF, England

*Abstract*

For any relationship to work there is a need for trust. The relationship between man and systems must evoke feelings of trust in the user otherwise the operation of the whole system (man with machine) will not be effective or efficient. Thus designers must ensure that systems are built in a way which allows the development of a trusting relationship. One area where this is of special significance is in the production of information for use in emergency situations. Here interactions must be swift and smooth; there will be times when there is no room for questioning of the information and instructions which the system is providing. This is particularly true in the world of aviation where action is often needed within a short period of time, where actions taken by the crew can be irreversible, and where there is little room for error.

One method of achieving trust in the system is to base design on user models of the system's operation and produce information and instructions that reflect the goals and expectations of the user. This paper outlines an approach to the design of a warning system for use on civil aircraft which incorporates this concept. The approach is based on models of system operations which are compared to the real world by model based reasoning and other AI techniques. The paper also discusses the current limitations of the model so far produced and the elements that are required to produce a more trustworthy system covering all aspects of the aircraft operation.

## 1. INTRODUCTION

Flying was initially an uncomfortable, difficult and hazardous experience. Early aviators soon found that the addition of a few basic aids helped them to achieve an acceptable and safe level of control, this heralded the development of aircraft instrumentation. As the functionality of aircraft instrumentation became more complex issues surrounding the interaction between user and system became more prominent. Current design guidelines indicate that the needs of the user must be met in terms of type and presentation of information provided, and that the design should allow for an interactive relationship to develop between user and machine. In order that this relationship operates effectively a level of trust in the system has to be arrived at and maintained. One way of seeing if this relationship has been working is to consider aircraft accidents and investigate the degree to which they can be attributed to a breakdown in this man-machine relationship and thus to see if system design is meeting requirements. An analysis of major accidents over the past ten years finds pilot error cited as the cause in 75% [1]. This explanation of "pilot error" is frequently used, but it does not always give the full story of how and why an accident occurred; it is just where the buck stops. However, it does indicate that at times the relationship may have broken down. It is in these high stress, abnormal situations, that trust in the system is most required especially as:

- **Time is limited.** Information and actions must be carried out in a timely manner and there is no room for error;

- **Multiple warnings can cause confusion.** There are two types of situation on an aircraft where multiple warnings are flagged simultaneously. The first is the genuine, although rare, case of simultaneous multiple failures. The second, more common case, is that of cascade warnings where failure in one system causes apparent failure in dependent systems;

- **More information is available.** Increased numbers and complexity of systems and increased measurement of system parameters has produced an increase in the amount of information available;

- **Aircraft systems are not independent.** The activity of one aircraft system is interlinked with other systems, therefore a change in one system is likely to lead to changes in other systems or flight parameters. These interrelations are currently not considered in most abnormal procedures and checklists.

This paper reports on the results of a project set up to investigate the application of Model-Based Reasoning (MBR) to warning systems with a view to providing trustworthy, integrated, user-orientated information about system malfunctions, abnormalities and failures.

## 2. AN INCREASING NEED FOR TRUST

On the flight deck there are both specific warnings, lights/lit legends or message directly related to an abnormal event; and indirect warnings, information such as the presentation of limits of single or multiple parameters. The development of warning and caution systems for aircraft is inextricably linked with the development of aircraft instruments, and with the growth of the capability of technology available to the aircraft and flight deck designer.

*Indirect Warnings and Historic Development*
Instruments on the Wright Flyer in 1903 were not fitted for the benefit of the pilot but for recording the performance of the machine for engineers to make a comparison with theoretical predictions. 1909 cine-film appears to show a Wright brother aircraft with the first instrument 'designed' to help the pilot, a piece of string in front of him, giving an indication of the angle at which the air was hitting the aircraft [2 & 3]; a warning of possible slip. Speed and attitude indicators were added to the flight deck by the first world war, but it was a technological break through, the usable gyroscope, that allowed the development of an artificial horizon which significantly moved instrumentation design forward. As instrumentation increased the need for the crew to develop trust in the information presented also increased, without this the information would not be used. An early example of this need for faith and trust in the instrumentation was demonstrated by James Dootlittle who flew on instruments alone in September 1929. When the automatic pilot, the first flight deck automation, was added in the 1930's the pilot had begun to share tasks with the machine and systems. As automation has increased pilot trust in the information provided, and the system performance has become even more significant.

During the 1950s a second major development occurred when electronic servo-driven instruments became possible. Now sensors could be placed remotely from the instrument, thus many new parameters could be brought onto the flight deck. Remote sourcing however removed the pilot's ability to check up on the system directly. Faith in instrument readings became increasingly necessary. Over the next few decades, systems design was consolidated until the third major technology change occurred; the introduction, in the 1980's, of what are known as integrated multi-function displays. These displays were initially based on Cathode Ray Tubes (CRT), but now there is a shift to Liquid Crystal Displays (LCD). Integrated displays have created significant changes in the opportunities for displaying and controlling information on the flight deck. For example the crew can now be presented with information sets tailored to their needs at any given time. However although these displays are known as "integrated" the information displayed is not truly integrated, formats consist of information elements displayed together on a common surface, it is not combined in anyway, the parameters shown are generally related to individual sources of data. This approach is designed to cut down scanning, and save of flight deck panel space, allowing the time sharing of high priority panel space [4].

*Aircraft Warning Systems - Direct Information*
As indicated above a great deal of information associated with warnings and diagnosis is available indirectly from instruments on the flight deck used primarily for normal flying tasks. Warnings themselves must get the crew attention quickly, irrespective of current eye position or workload and must facilitate the right corrective action quickly. Specific warning systems began initially with a fire bell and a couple of warning lights. As the number and complexity of systems on the aircraft have increased and the ability to measure system parameters has improved, so the requirement for alarms and warnings has also increased. e.g. from the Boeing 707, 188 warnings to the Boeing 747 with 455. Early aircraft had numerous warning lights spread around the flight deck (e.g. Buccaneer) [5], with increased numbers of systems and associated warnings, the central/master warning/caution feature was developed. This groups together warnings in the central field of view for both pilots. These system [6] were initially a group of lights mounted in the glareshield, some of which have an associated audio warning. These grouped warning functions together by system e.g. electrics. The Lightning was the first military aircraft to have this feature [5]. Specific coding for warnings was also developed using colour to represent urgency; red for those requiring immediate action; amber for those requiring immediate attention but not necessarily immediate actions. However, on most aircraft these are only single legends and additional information is required from other sources such as other instruments, manuals, and aircraft behaviour to diagnose the situation.

These technology developments have provoked and accompanied a change in operational philosophies.

The move to multi-function displays enabled all information to be provided to the flight crew enabling the move to two crew operations. the increased reliability of aircraft systems invited the introduction of a 'need to know' approach to the provision of system information [7] first used in the A300 in the 1970's. The main elements of the modern approach to warning systems are;

- **Quiet & Dark.** The absence of visual indications of normal conditions;

- **Simplification.** Simplification of failure information (systems displays) and presentation of corrective actions (warning displays);

- **Decreasing Discretes.** Reduction in the number of different attensons, using a basic attention getter to alert the crew to the presence of warning messages;

- **Phase of Flight Sensitivity.** Warning displays adapted for the requirements of different phase of flight, with suppression of warnings for safety reasons at the most critical periods (e.g.. take-off and landing).

The A320 is probably the most advanced system in use [6 ] which employs this philosophy. It provides primary engine instrumentation, the warning and alerting functions, together with synoptic displays and checklists of necessary actions on two CRT displays.

*Limitations*
There are, however, unforeseen side-effects of this development process. Hand in hand with the increased automation and reduced crew size, has come a decrease in direct access to information about the status and activity of the aircraft systems. There is a loss of constant availability of many sensed readings, which can provide trend information that could be used to predict problems prior to their onset. Many systems now only call the crew's attention when parameters pass out of limits i.e. beyond pre-determined fixed thresholds. There is increasing concern that in unusual circumstances the crew no longer have the information, experience and skill to solve the problems that may arise. It is the crew's capacity to analyse, seek novel solutions and extrapolate beyond the immediate situation that is required.

Accident/incident investigations have shown that automated systems can adversely affect crew communications and decision making as a result of the crew's reliance on the systems fostered by the reduced requirement for systems monitoring. The

crew must now have faith in their system's ability. This faith is a fundamental element in the construction of a trusting relationship. The techniques applied to central warning systems which this paper describes have been used with a view to instilling faith and trust in the information and advice which the system provides.

## 3. BUILDING TRUST THROUGH USER ORIENTED SYSTEMS DESIGN

One indication that a relationship between the system and the user is sound is that the user is able to follow implicitly the reasoning of the system. If the underlying logic of the system is alien to that of the user misunderstandings, breakdowns in communications etc. will occur, and the user's confidence and trust in the system will be eroded. Artificial Intelligence techniques such as expert systems are intended to replicate human reasoning about problems.

In recent years research has been carried out on expert systems to aid the performance of the operator in fault finding and diagnosis. The resulting systems might advise unskilled operators; help overloaded skilled operators; or perform some functions automatically where operators are highly overloaded. Many of the systems produced have been based on rules which are matched to a set of conditions. However, this approach means that the ways in which the system can fail and the symptoms associated with that failure are pre-determined Thus the system is apriori in nature [8] and any symptoms or failure modes which are not included within the rule set at system design can not be handled. This apriori approach would not be suitable to the warning system application because of complexity of the interaction of the various aircraft systems creates a considerable number of failure modes and an even greater number of possible indications of incipient failures. An alternative approach, namely model-based reasoning (MBR) has been adopted as it operates on an 'understanding', or model, of the underlying system; an understanding which is built on an idea of how the individual systems work, and how they interrelate.

In general terms the MBR approach uses models to represent and 'understand' situations, events and systems. A set of models monitor system response at a local level, at the lowest level receiving information from the aircraft systems assessing the status and trends of the output parameters of the aircraft system it represents. The purpose of these models is to monitor the health, status and configuration of the associated aircraft systems,

reporting any faults or abnormalities in behaviour. The models are interconnected in a network with the links representing the interaction between system components (the model network). The nodes report system status to an executive function, the Network Executive, which converts the information it is receiving forming a view of the overall situation. From this view the executive forms a plan, an ordered set of actions; derived from the standard operating procedures to ameliorate the effects of abnormalities, rectify the effects of malfunctions and maintain the safe operation of the aircraft.

The models within the system use parameter values which are available to the crew and the actions which can be recommended are based upon Standard Operating Procedures and are therefore the same as those the crew would take if they had sufficient time to diagnose the problem and compose the action list. The crew and the system therefore have access to the same information and are constrained by the same operating rules; the system behaves as though it had the same training as the crew.

### 3.1    Matching Models To User Understanding

The models of system operation used in this approach are based on details of the aircraft systems. The level of detail within the models was constrained in two ways in order to match their operation and the data provided to the needs of the user and the avionics environment. The first constraint is that when a system fails or gives problems in-flight the ways in which aircrew can respond are restricted by their ability to access the systems. Normally they can either reduce usage of the system; use an alternative system to achieve the same functionality or lose the functionality [9]; most problems can not be fixed until the aircraft is on the ground. Therefore only certain information is of use to the crew. The models used have been designed to meet the needs of the crew and expanded a little to provide the additional information required to achieve a degree of diagnostic functionality that could be used on the ground. The second constraint on the level of detail necessary within a model is related to the level of information available on the aircraft. A model can only use and monitor parameter values that it has access to thereby reducing the complexity necessary. However, the aircrew will have access to the same range and detail of parameters and so the model is constrained to interpret that data which is available to the aircrew.

This match of the models to the needs of the user and what they can achieve in the operation of their task and the use by the models of data available to

the user are a design features which can help to build trust in the system's capability. Provision of information that did not meet the crew's needs would tend to increase workload and undermine any relationship. Also presenting information at a level of detail that is beyond the needs of the user or that is of no use to the user will increase workload, serve to confuse, and give the user a poor view of the system's ability to make a positive contribution, again jeopardising any trust.

### Operation of the Models
With these constraints in mind the modelling technique which appeared most appropriate to the problem was the use of fuzzy logic models as used in fuzzy logic controllers currently being evaluated for a wide range of applications in which human expertise is applied to optimise the control function [9].

At the core of such a model is a description of the correct operation of the system being modelled rather than a list of the ways in which the system can fail. Any discrepancy between expected correct operation and observed operation is then a fault or a malfunction. The description of the system being modelled can be derived from a number of sources including operator experience, design definition etc.

As an example consider the model of a fuel pump in Figure 1. This model is being used to monitor the health of the pump without having any direct sensing of the pump. In this application interest in the health of the pump is confined to an assessment of the quality of its operation i.e. its fuel flow delivery pressure. A simple descriptive model of the pump would be:

IF there is sufficient fuel in the tank AND the pump is switched on AND the pump is connected to a powered bus THEN the pump pressure will be NORMAL OTHERWISE it will be LOW.

This model contains a number of descriptive elements. Those concerning the fuel level, the switch position and the power supply can be considered as binary. The others, NORMAL and LOW, applied to the pump pressure would be adequate and usable for the human operator but are not amenable to machine manipulation unless they can be defined more rigorously. The use of fuzzy sets allows this definition.

Essentially the use of fuzzy sets is the mapping of a numerical value onto a set distribution. In the example of pump pressure three sets are considered suitable for a description of the value. The sets used

are illustrated in Figure 2 where a value can have a degree of membership between 0 and 1 of any two sets. Thus a pump pressure value can be represented as a vector (x, y, z) where x, y, and z all lie between 0 and 1. When the pump pressure is NORMAL then the vector will read (0, 1, 0) and when the pump pressure is LOW the vector will read (1, 0 , 0). The model then states that:

IF there is sufficient fuel in the tank AND the pump is switched on AND the pump is connected to a powered bus THEN the pump pressure is (0, 1, 0) OTHERWISE it is (1, 0, 0).

The sensed value of the actual pump pressure is then mapped onto these sets and comparison is made between the expected i.e. model, distribution and the actual distribution. Mismatches in this comparison process provides a fault detection mechanism.

The threshold value used to give the conventional low pressure warning is set at the point in the distribution where the vector takes the value (1, 0, 0). It is therefore possible to detect an incipient malfunction in the pump's performance and monitor the progress of the malfunction, is it deteriorating, improving or stable, before failure occurs. This can be used to attract attention to an emerging situation and thereby provide additional time for crew response [10].

The interlinking of a large set of these simple models provides the means of fault isolation i.e. identifying where the fault has occurred rather than where the symptoms become evident. In the example used above the pump model required a number of conditions to be met in order that the delivery pressure would be within the NORMAL range. These conditions included fuel level in the relevant tank and power supply from the relevant electrical bus. This information would be derived from models of the fuel tank and power bus. Thus the pump model would determine that, if there were no power coming from the bus to which it is connected then the pump pressure should be LOW. If the actual pump pressure is found to be LOW the system will consider the pump to be operating properly and that the loss of delivery pressure is caused by the lack of electrical power on the bus.

This process of fault isolation allows the system to identify the primary causes of a failure and avoids the problems presented to the user when indications of malfunction cascade through the aircraft systems.

This approach appears to be readily understandable and cognitively accessible to the user. Use of fuzzy

logic allows the system to reason in a manner that reflects that utilised by humans and provides information about trends and relative values. Use of knowledge of human information processing as incorporated in these models, leads to the design of systems that will reason about problems in a manner similar to that of the operator and therefore will be understandable and aid the generation of trust.

## 3.2 Providing Output That The User Will Use

### *User Expectations and Goals*
An indication that the relationship between the system and the user is working is whether or not the user will following instructions given by the system. If the outputs are not, or do not appear to be, correct the user will not undertake them, possibly leaving the aircraft in an unsafe configuration. The perceived correctness of the instructions is dependent upon an alignment between the user's and the system's understandings of the situation. The provision of trend information as described above enables the user and system to share a common view of the situation. Also the information provided must be sufficient for the user to understand the situation, but must not be at a level of detail that will serve to confuse or introduce work overload. This is achieved by various design features implemented in the system to assist the user.

The configuration of the aircraft could be described by the position of individual switches controlling the systems and failure states. It is this configuration that has to be controlled in response to the information provided by the warning system. The purpose of the control task is to maximise the safety of the aircraft in abnormal failure conditions. The position of any one switch or state of any one system cannot be viewed in isolation since each makes a contribution to the overall configuration of the aircraft, and one system state can influence the state of another. The actions advised by a warning system such as the model-based reasoning system described here are concerned with reconfiguring the aircraft systems.

Currently when single system failures occur the crew relies on the procedures laid down in the aircraft's Quick Reference Handbook for the optimum solution. However, these procedures do not cover multiple failures and unusual configurations. In these cases use of them could drive the crew to place the aircraft into an unsafe configuration, for several reasons:

- **Contradictory Actions.** Some of the actions specified by single procedures to solve different

elements of the overall problem may be contradictory;

- **Combined Failures.** Two required procedures may specify actions which when combined remove essential facilities;

- **Consequential Failures.** Loss of system functionality can occur because a supplying system is switched off. This form of failure is difficult to identify as consequences of system changes can spread throughout the aircraft, it is particularly hard to anticipate when systems are in unusual configurations due to other failures.

When the crew are confronted with these problems they must use their knowledge of the aircraft, the limited procedure information available and their understanding of the situation to invent solutions which are compatible with their operating procedures. Coping with these type of problem is fundamental to the operation of the Network Executive.

In single failure situations the Network Executive is required to produce the standard procedure as set out in the Quick Reference Handbook. However, when multiple failures occur the isolated fault procedures may not be the optimum solution. The Network Executive uses the concepts of a goal hierarchy and constraint satisfaction methods to generate sequences of recommended actions. These are managed and integrated into user checklists for complex or multiple failure situation.

The crew have a number of goals to meet which arise from their task to get the aircraft to its planned destination, ensure the safety of the operation and maintain passenger comfort. Each of the goals will have different priorities that may vary with phase of flight and other events. Also goals are rarely independent of each other thus decisions about one involve the consideration of others. For example, switching off an engine will involve the loss of any generator run by that engine and consequently, at least temporarily, any systems run by the electrics associated with that generator. The crew must take the information that is available, their knowledge and experience and decide on the best course of action. The best action being that which will achieve the highest number of goal 'points' possible, (given that higher priority would have a higher 'score' associated with it); e.g. if the engine is on fire loosing the generator to make the situation safe would be small price to pay.

The system must understand the interrelationships and the values associated with different actions and outcomes to provide the aircrew with the best advice and suggested actions that will optimise the solution. If the system cannot achieve this and suggests apparently inappropriate actions to the user it will appear not to understand the interrelationship of systems and the consequences of the actions. Hence the user will not have faith in the systems ability to assist in decision making process. Under these circumstances trust would break down and efficient use of the man-machine system would cease. The advice provided by the warning system must therefore be orientated towards achieving the user's goals.

Keeney and Raiffa [13] advocated an approach that was initially thought to be a possible solution to this issue. In their theory each of the goals of a system (or objectives) is given a weight value which can be used as a measure of the success of a particular solution. However there are difficulties with this approach within this application since the weights must be obtained through extensive interviews with experts in the total system. The difficulties of detailed knowledge elicitation is notoriously, and it is difficult to find an expert with a sufficiently global view covering the aircraft operations, the complexity of failures and the interaction of aircraft systems. A more simplistic approach is to order the objectives and use fixed lexicographical ordering of possible solutions. Such an ordering is much easier to develop and the solutions reached by it are easily explicable to the crew who can appreciate the reasoning, even if they disagree with it.

Solving the problem of multiple levels of goals alone will not solve the total problem of configuration management, because of the interrelated nature of the action of the systems. There are multiple ways of satisfying each of the goals, therefore some form of representation of states that could achieve the goals is required, checking that all the required goals have been achieved. The approach selected for the design of this system was based on Goal Tree-Search Trees (GTST) [14]. This approach involves the decomposition of goals into subgoals and the equipment states that they rely on. Each goal is functionally decomposed into contributing states to build a rule-based tree ultimately linking equipment states to goal states by many intermediate levels. The tree structure can be based on the knowledge of domain specialists or from the systems design rather than from a global expert. Unlike rule-based systems, which require heuristics and underlying rules from experts, the knowledge required is easily specified and obtainable. The nature of aircraft systems means that the resulting tree is intertwined, a network which makes the search for optimum

solutions difficult. The network is in fact a form of model-based system itself. It models how the systems states contribute to the high level gaols, not to the individual systems.

Using this approach the Network Executive also prioritises the actions required and separates them into logical groups. This is achieved through a multiple objective representation of the problem. These objectives, or goals, are obtained from discussion with pilots and analysis of the aircraft systems, and are prioritised in order of importance. When a failure occurs the optimum solution is calculated, any system state changes identified are related to the goal that requires them and prioritised in accordance with the goal's position in the fixed hierarchy of objectives. Thus what appears to be a complex single procedure is broken into several meaningful sub-procedures related to individual objectives thereby aiding understanding and situational awareness. This checklist information can be grouped under 'why' an action is required, say to stabilise an engine or reconfigure electrics; giving the user further information about what is happening and why the actions are recommended. .

*Text for Instructions and Integrated Checklists*
The system must produce an output which covers all failures and abnormal situations and produce optimal corrective actions across the aircraft configuration. After the identification of actions required, outlined above, the next stage is to produce a list of actions for the user. User's require actions which are familiar, which they can do, and which must take account of the whole aircraft configuration. Information presented to the crew must be logically ordered from a system and user viewpoints. The reasoning behind the information provided must be implicit, so that the user can follow what is going on. If this is not included in the design the user will not have faith in the ability of the warning system to meet their needs and the system will not be used. The MBR approach tackles this issue by the use of the goal hierarchy and procedural elements to solve each of the individual elements of the control of systems configuration. The warning system contains a set of minimal actions, e.g. switch off, in an actions table. These elements can be constructed in language, and a dialogue structure similar (if not identical to ) the individual checklist items currently available on the flight deck. This means that instructions will be familiar and the meaning of these elements will be consistent at all times. This consistency and familiarisation allows the user to build up trust and faith in the system.

For single failure situations these checklist would look, and contain actions, identical to those in the Quick Reference Handbook. When multiple failures occur the new procedures will be required. Each individual procedure generated is associated with a goal in the network structure. As these goals are prioritised in order of importance the procedures can be prioritised under the same mechanism. These are then grouped under logical headings of actions to attain a goal rather than the current systems orientated which does not work in the multiple failure case. Examples of these goals are Making Engine Safe or Resetting The Electrics. If the crew take action other than that advised the system will continue to generate advisory actions based on the current configuration of the aircraft

## 3.3 Additional Features

To operate as effectively as possible the user needs a complete picture of the situation, including the final status of the aircraft and its capabilities after a reconfiguration event. Providing this information will enhance the users view of the warning system's capability and ability to help them. The MBR system aims to achieve this complete picture with the capability to provide a variety of information not provided by current systems.

*Consequence Anticipation*
When a failure occurs the algorithm searches for an optimum solution, this solution does not only include the procedure to deal with the warning, but also advises of all other actions necessary to achieve the optimum state. These secondary actions will be responses to the effects of applying the initial procedure. For instance, when an engine fire occurs the engine must be switched off. In order that the electrical supply is maintained it may be advisable to switch on the APU and re-route the electrical supply by closing various bus ties. These actions are anticipated and presented before the initial action of switching off the engine is taken. The provision of advice in this manner allows the crew to see that the system has considered the impact of the recommended actions and that it is responding to the situation intelligently i.e. as the crew would if they had sufficient time to consider all options. This will increase the level of trust which the crew will place in the system.

*False Warnings*
In addition to providing correct data the system must not provide incorrect data. The production of false or nuisance warnings is a feature of a system that will rapidly undermined user trust. Discussing the development of Airbus ECAM approach Potocki De Montalk [11] sees timely and appropriate alerts as a

major factor in successful design. Guidance as to actions to be taken and information about how the situation will affect subsequent operations are important. Reliability and integrity are increasingly important factors in the design of warning systems. There are problems with both unwanted alerts and failures to alert [12]. Warning systems must eliminate events in the areas B and D in Figure 1. The fault detection and isolation process described is independent of fixed thresholds and thus has the potential to detect sensor failures and help eliminate numbers of false warnings.

## 4. FUTURE DIRECTIONS
*Aircraft Operations*
The current system has no knowledge of the aircraft operational goals; its destination; operational constraints etc. The advice which the system produces needs to be modified by the crew's use of this knowledge For example, if the aircraft has a slow leak in one of the hydraulic systems the impact of this will differ depending on the position of the aircraft within its operational plan. If the aircraft is flying long distance the fault will have more of an impact on the actions required than if it is only flying a short distance or it is close to its destination when it occurs. Much of this information may be available to the system through the Flight Management System and the navigation systems.

*Dynamic Goal Hierarchy*
The current system works with a fixed hierarchy of goals, however the ordering should be dynamically based on the different phases of flight. In future work these different hierarchies will be identified and the transition between them managed.

## 5. CONCULSION/SUMMARY

The operational features required of the warning system to met the needs of the user in abnormal conditions on the flight deck are achieved by the model-based approach. The features of this approach are:

- **Early detection.** Detection of abnormal condition development before action is required, giving early warning, enhancing situational awareness and maintaining a common understanding between system and crew;

- 

- **Isolation of primary cause(s) of failure conditions.** Sorting information into a flight related priority list for action, helping the crews to effectively process the larger quantity of

information effectively and inviting the crew to trust the system;

- **Ordering actions.** Provision of ordered action checklists in the event of single and multiple failures based upon standard procedures giving the appearance of common training. This helps direct the crew to the correct actions, even if the individual actions for multiple failures would conflict, thereby assisting the crew in novel situations.

- **Indication of consequences of actions and inaction.** Currently the crew are provided with little or no information about the consequences of their actions or inaction. In some cases actions once taken cannot be reversed; knowledge of this in advance may effect the most appropriate actions, especially in multiple failure situations. This encourages confidence that the system knows what it is doing.

- **Awareness.** Continuous awareness of system's configuration, and enhanced awareness of health and status of aircraft systems. Situational awareness for the crew, the ground (for planning) and maintenance.

The quality of the information provided can be increased and a potential reduction in crew workload could be achieved by the introduction of the MBR approach to warning systems. This in turn would lead to an improvement in the quality of decision making. This would increase operational effectiveness and safety. This will, for example, be through the system's;
- ability to automatically prioritise warnings in multiple failure situations, in relation to the needs of the flight.
- giving advanced warning of potential problems allowing for greater planning and more possible options to be available.
- providing advice on decreasing the use of a system to maintenance capability for more critical phases of the flight.
- reducing nuisance and false warnings. This is achieved by the application of reasoning logic to the situation. This will decrease workload and reduce the chance of the crews being mislead by them.

The design approach adopted is user-orientated employing models which replicate the user's model of the aircraft system and which operate on the same data that is available to the user. The advice which the system provides is built up from the procedures which the user has been trained to undertake. the system should, and has, provoked the response, "It

thinks like I do", and thereby encourage confidence and trust in the information and advice which it provides. With the production of the capability at the user interface to develop a man-machine relationship that has its foundation in trust the man-machine system can perform more effectively.

## 6. ACKNOWLEDGEMENTS

## 7. REFERENCES

1. Weener, E., "Boeing Blames The Pilots For Disasters.", Guardian 14 November 1990.
2. Chorley, R. A., "Seventy Years of Flight Instruments and Displays", Journal of the Royal Aeronautical Society.
3. Chorley, R. A., "Trends in Flight Deck Displays" Smiths Industries TSP 2674, 15 Sept 1977.
4. Pischkle, K. M., "Cockpit Integration and Automation; The Avionics Challenge", International Federation of Airworthiness International Conference, Nov 1990.
5. Gorden-Johnson, P., "Aircraft Warning Systems: Help or Hindrance Philosophy of Warning Systems", IMECH E 26 March 1991.
6. Alder, M. G., "Warning Systems for Aircraft: A Pilots View", Philosophy of Warning Systems, IMECH E 26 March 1991.
7. Speyer, J-J., Monteil, C., Blomberg, R. D., and Fouillot, J. P. ,"Impact of New Technology in Operational Interface. " In design Aims to Flight Evaluation and Measurement", AGARD-AG-301 Vol. 1, March 1990.
8. Frankish, C. R., Ovenden, C. R., Starr, A. F., and Noyes, J. M. "Application of Model-Based Reasoning to Warning and Diagnostic Systems for Civil Aircraft. ERA Conference, Affordable Technology in the Civil Military and Space Sectors - Opportunities and Challenges", December 1991.
9. Ovenden, C. R., "The Structural Model Network", Smiths Industries, RPT, MBR 15/1, December 1993.
10. Ovenden, C. R., "Model-Based Reasoning Applied To Cockpit Warning Systems", Advanced Aircraft Interfaces: The Man Side of the Man-Machine Interface, 1992, AGARD-CP-82.
11. Potocki De Montalk, J. P. "Airline Warning/Caution and Maintenance Systems. Philosophy of Warning Systems", IMECH E, 26 March 1991.
12. Newman, T. P. "Chephalic Indications, Nictations and Anopic Ungulates" Philosophy of Warning Systems" IMECH E, 26 March 1991.
13. Keeney, R. L., and Raiffa, H. "Decisions with Multiple Objectives: Preference and Value Trade-offs" John Wiley and Sons.
14. Kim, I. S. and Modarres, M. ,"Application of Goal Tree-Success Tree Model As The Knowledge Base of Operator Advisory Systems", Nuclear Engineering Despatches, 104, 67-81.

## 8. FIGURES



Figure 1: Fuel Pump Model



Figure 2: Pump Pressure Distribution



Figure 3: Integrity of Warning Systems; after Newman (1991)

# Trust and Adaptation Failure: An Experimental Study of Unco-operation Awareness

Taylor, R.M., Shadrake, R. and Haugh, J.
DRA Centre for Human Sciences, Farnborough, UK.

## 1. SUMMARY

The effects of variations in the adaptation performance of automation were investigated using a computer simulation of aircraft tasks, with levels of automation aiding under experimental control. Failures were introduced in the ability of the automation to intervene in a timely and appropriate manner, with regard to the prevailing tasks demands. Measures of performance and subjective assessments were recorded. Compensation for performance occurred without awareness of automation failure. The implications for the design of safeguards against adaptation failure are discussed.

## 2. INTRODUCTION

Building safeguards against automation failure is a primary concern for designers of highly automated, mission critical, and safety critical aviation systems. While totally fail-safe flight systems are impossible to engineer, the designer must rely on operator intervention and manual control as default to maintain safety, with appropriate margins for human error. In advanced automated systems employing Human-Electronic Crew (H-EC) concepts such as teamwork, co-operative functioning, and adaptive aiding, functions seem likely to be shared increasingly between human operators and computers. In such H-EC systems, trust seems likely to be the *psychological glue* which holds together the functioning of the human system components. However, over-trust can adversely affect reactiveness to automation failure. Also, loss of trust, arising from unreliable or inaccurate machine or computer performance, can be the cause of operator unwillingness to utilise automation functions. Such reluctance, if unwarranted or misplaced, can lead to inefficient system operation and ineffective human-system performance. Thus, safeguards are needed against automation failure that ensure that operator trust in automation functioning is maintained at realistically appropriate levels. Sources of automation failures can arise from breakdowns at different levels of system functioning. Advanced H-EC systems involve increasingly high levels of human-computer functional integration and interaction. There is a need to develop understanding of how humans might react to, and cope with, high level functional failures, and to develop appropriate high level functional safeguards.

### 2.1 Trust

Concern about under-reliance on automation, and undesirable manual intervention has led to interest in the social dynamics of trust between humans and machines. Investigations of the quality of teamwork in RAF aircraft tactical missions show that teamwork performance is affected by trust (1). Trust was a significant factor in distinguishing between good and poor teamwork performance. Trust was rated at a significantly lower level in single-seat RAF Harrier operations (i.e. human-computer teamwork) than in two-seat RAF Tornado aircraft tactical operations (i.e. both human-human and human-computer teamwork). Experimental evidence has verified that unexpected automation failure leads to a breakdown of trust, and to difficulty in the recovery of trust with a loss of faith in future teamwork performance (2,3). As trust declines, manual intervention increases. Other research has investigated how when workload is increased, over-trust or complacency develops with automatic systems (4). Complacency, coupled with vigilance problems, is likely to lead to failure to detect performance deviations and decrements in automation performance. People generally distrust computers. Trust between humans is engendered by continuous, repetitive, and reciprocating actions. In the same way, it is plausible that H-EC trust will build-up when automation performance conforms consistently and predictably to expectations, in accordance with agreed mission goals.

Understanding the factors that affect trust could help design safeguards. At the 1st HE-C Workshop in 1988, an investigation was reported of trust in two-seat RAF Tornado aircraft tactical operations (5). Tactical decision-making scenarios were elicited and rated for the importance of factors associated with trust in the events described. These subjective ratings showed that the demand for trust was associated with the perceived risk and the probability of negative consequences, whereas the supply of trust was related to the requirement for judgement and awareness, and the uncertainty and doubt in making the decisions. Thus, relying on others to make risky decisions calls for a large amount of trust. But if the decision requires another person exercising a high degree of awareness and judgement, and there is much uncertainty and doubt in the decision provided, then the actual trust engineered by the decision will be low. At the 2nd HE-C Workshop in 1990, a model was proposed of the relationships between trust, operator skill level, task complexity, workload, risk, self-confidence, and EC reliability (6). Subsequent studies, in which workload and automation reliability were varied, led to refinement of the model to include the factors of fatigue and learning about system states (7). Other research has shown how trust can be modelled as a function of parameters such as recent performance, and the presence and magnitude of a fault (8). Intervention and automation use are influenced by the combination of trust and self-confidence in operators' abilities to perform the task by manual control. Operators will allow automation to have control if they trust it; and they will take control themselves if they distrust it, providing self-confidence is sufficiently great. However, high self-confidence often produces a bias in favour of manual control.

### 2.2 Adaptive Aiding

The concept of adaptive aiding introduces complex demands for trust of computers. Adaptive automation is an approach to human-systems automation in which the control of the onset, offset and form of automation for a specific task or function is mutually shared between the human and the adaptive automation systems. Responsibility for performance or execution of tasks can

be solely that of the human or the machine, or can be shared between those entities depending on the particular strategy of automation used. In a recent review, Rouse (9) describes how early experiments with an artificially intelligent decision aiding system prompted conflicting intelligence, where the human and computer independently chose reasonable but conflicting courses of action. Co-operative rather than conflicting intelligence was needed so that the human and computer could complement each other to yield improved performance and safety. This human-centred approach was predicated by the condition that the default between automatic and manual control should be manual. Humans should be able to perform any tasks they choose to perform; automation or aiding should be available in the event that humans choose to delegate authority to perform tasks. Human judgement should be the ultimate authority and the operator should be able to take charge. The idea that computers could take tasks, but not give them to humans, is called the *First Law of Adaptive Aiding*.

In considering how to create co-operative intelligence, it became apparent that it was inefficient to strictly allocate functions and tasks to one performer. A queuing model of human-computer interaction predicted the advantages of a flexible, dynamic allocation of tasks among two servers over a strict, static allocation. System initiated automation, with either explicit or implied pilot consent, could help in situations of unexpected high workload. It followed that the nature of computer assistance should vary with the situation, including the state of the tasks, the environment, and the humans in the system. From this emerged the concept of adaptive aiding, with allocation, partitioning, and transformation of tasks in response to current and impending situations.

## 2.3 Automation Risks

Adaptive automation introduces new risks for successful system functioning, and along with it, the need for safeguards. Dynamic task and function allocation, with a manual default allows the possibility of unnecessary and inefficient manual intervention. With adaptive automation, different roles and responsibilities may be assigned for the same tasks at different times depending on the particular automation strategy being invoked. But this variability could easily lead to an appearance of total unpredictability unless care is taken in the design and implementation of adaptive automation. Dynamic allocation of tasks only makes sense *if all the performers are aware of what each other is doing* (i.e. both human and computer task awareness). Otherwise, tasks might be overlooked or task contention might occur.

Sharing tasks and functions between humans and the computer introduces the risk of over-reliance and dependency on computer aiding. Reliance and dependency lead to reduced system awareness and degradation of manual skills. This becomes a problem in the event of automation failure requiring manual intervention as the default. In dynamic systems, when the information relevant for decision making changes over time, and is not static, a dynamic internal model of the task is needed to guide decision-making. An appropriate dynamic internal model of the important changing relationships will be difficult to maintain for

regaining manual control following automation (10). Another problem is that operator detection of automation failure is degraded with a static allocation fixed over a period of time (11). Monitoring automation performance for failures is inefficient, due to a natural tendency towards complacency, and because of difficulty maintaining vigilance without active task involvement.

Because of problems with failure detection and manual skill degradation, manual task reallocation has been proposed as a countermeasure to monitoring inefficiency and complacency. It has been shown that short periods of intermittent manual task reallocation, or cycling between manual and automation control, reduces failures of monitoring (11). By maintaining manual skill levels, and enhancing situational awareness, manual task re-allocation helps in the event that intervention is needed following automation failure. However, automatic re-allocation of tasks to manual seems close to a violation of the First Law of Adaptive Aiding, that the computer should not give tasks to humans. Also, as referred to earlier, without careful consideration of the procedures needed for implementation of dynamic task allocation and re-allocation, such variable assistance and allocation could lead to unacceptable unpredictability.

## 2.4 Issues

The following questions arise from consideration of the possibility of adaptation failure. They are indicative of areas of uncertainty for further research:

a. *Trust.* What happens when the computer fails to intervene appropriately, or if there is an automatic reallocation to manual during a period of high workload? How does variability in the adaptive strategy affect trust and other attitudes that are likely to affect automation use? To the extent that trust is an attitude, does misplaced trust create cognitive dissonance, and if so, how is the dissonance resolved? If trust in adaptive functioning is lost, can the appropriate level of trust be recovered?

c. *Awareness.* Awareness of task performance is important for dynamic function allocation. In human teams, there can be awareness of other individuals performance, i.e. someone not pulling their weight. In working with an adaptive co-operative system, variations in task demands will arise from the external situation, and from the EC's contribution to the team's performance. Is awareness of adaptation failure, or of changes in the adaptive strategy, a system design problem?

d. *Compensation.* Human teams naturally compensate for variations in the performance of individual team members. Can the operator compensate for the computer by picking up the extra work? Is awareness of the variations in automation performance necessary for compensation to occur? If the same compensation were to occur bidirectionally in the H-EC team, could this be one mechanism which might engender trust in the team by others outside the team?

## 3. EXPERIMENT

### 3.1 Aims of Experiment

An experiment was conducted to investigate the effects of variations in automation adaptation performance on operator performance, and on attitudes that might affect automation use. In particular, the intention was to

discover the sensitivity of task performance measures, and of subjective rating scale dimensions, to automation adaptation failure, and to examine the structure of the relationships between these dependent variables.

## 3.2 Experimental Method

a. *Subjects.* Twelve non-aircrew subjects participated in the experiment. All were staff at DRA CHS.

b. *Task.* The Multi Attribute Task (MAT) battery developed by Comstock and Arnegard (12) provided the task environment. The MAT battery comprises a computer simulation of three tasks, namely tracking, monitoring and resource management, presented simultaneously. A diagram of the screen format is shown in Figure 1.



**Figure 1. MAT Screen Format**

The tracking task was a two dimensional compensatory tracking task that required subjects to keep a target in the centre of the tracking window. The monitoring task required subjects to monitor and correct deviations on four gauges. The resource management task was complex, and required subjects to maintain the fuel level of two main tanks at a specified level by transferring fuel from several supply tanks. All three tasks could be operated manually, via a joy stick and keyboard, and both the system monitoring and resource management tasks also could be operated either aided or fully automatic. In the aided mode, parts of the tasks were automated, leaving the subjects to monitor the automation performance and to complete the tasks. The tracking task was always manually operated.

c. *Scenarios.* On the experimental trials, subjects were presented with four scenarios. The independent variable was the manipulation of co-operation given by the adaptive aiding. In all four scenarios, the invocation of the automation occurred automatically, i.e. changes in the automation level were initiated by the computer, and not by the operator/subject. Co-operation was experimentally manipulated by providing two co-operative and two unco-operative scenarios. In all four scenarios, the frequency of events requiring action, and the resultant task demands, increased as the scenarios progressed. In the co-operative scenarios, the level of aiding provided by the automation increased appropriately with the event frequency and level of task demand. In one co-operative scenario, the system monitoring task went aided and remained so; the resource management went aided, and then fully automated (*Co-op 1*). In a second co-operative scenario, the resource management task went aided, and

remained so; the system monitoring went aided, and then fully automated (*Co-op 2*). In the unco-operative scenarios, the level of aiding initially increased appropriately, but then shortly after the onset of a period of particularly high event frequency and task demand, the aiding automatically re-allocated to manual. In one unco-operative scenario, the system monitoring task switched to, and remained at aided; the resource management task also went aided, then warned to go fully automated, but then reverted to manual (*Unco-op 1*). In the other unco-operative scenario, the resource management task switched to, and remained at aided; the system monitoring task also went aided, then warned to go fully automated, but then reverted to manual (*Unco-op 2*). Each experimental scenario was of five minutes duration. The sequence of events is illustrated in Figure 2.



**Figure 2. Event Sequence**

d. *Design.* All the subjects were presented with the four experimental scenarios. Subjects received a 5 minute practice session before each experimental session, which included combinations of automation and manual operation of the three tasks. All subjects received the same amount of practice on the task, including manual, aided and fully automated modes. Subjects were instructed that the computer would attempt to give the appropriate level of aiding, but that accuracy of the computer's judgement was imperfect and that compensation may be required if the aiding failed to be appropriate. The order in which the scenarios were presented was balanced to prevent order / sequence effects, using the four sequences described below:

| Subjects | Order of Scenarios | | | |
|----------|--------|--------|--------|--------|
| 1, 5, 9 | Unco-op1 | Co-op1 | Co-op2 | Unco-op2 |
| 2,6,10 | Co-op1 | Unco-op2 | Unco-op1 | Co-op2 |
| 3,7,11 | Co-op2 | Unco-op1 | Unco-op2 | Co-op1 |
| 4,8,12 | Unco-op2 | Co-op2 | Co-op1 | Unco-op1 |

e. *Dependent Variables.* Dependent variables comprised computer measures of task performance and subjective ratings. The following MAT task variables were recorded: root mean square (RMS) error on the tracking task; the number of correct resets, incorrect resets, and mean reset time on the system monitoring task; tank 1 deviations, tank 2 deviations, and the number of pump activations on the resources management task. After each scenario, subjects provided subjective assessments on 7-point Likert rating scales of the timeliness and appropriateness of the computer aiding on dimensions (low to high) of 17 constructs related to trust and awareness. Situational awareness (SA) was rated using 3-d SART dimensions (13).The constructs for ratings were defined as follows:

*Confidence* - Confidence in own ability to successfully complete the tasks with the aid of the adaptive automation; *Self Confidence* - Confidence in own ability to successfully complete the tasks; *Accuracy* - Accuracy of own performance on the tasks with the aid of the adaptive automation; *Self Accuracy* - Accuracy of own performance on tasks;

*Automation Confidence*-Confidence in ability of the machine to support successful completion of the tasks; *Automation Accuracy* - Accuracy of machine in supporting successful completion of tasks; *Automation Dependability* - To what extent can you count on the machine to provide the appropriate support to the tasks; *Automation Reliability* - To what extent can you rely on the machine to consistently support the tasks; *Automation Predictability* - Extent to which you can anticipate and expect the machine to support the tasks; *Risk* - The probability of negative consequences of relying on the machine to support successful completion of the tasks; *Impact/Survivability* - The severity and criticality of adverse or negative consequences of relying on the machine to support successful completion of the tasks; *Decision Complexity* - The extent to which the machines' decision on when and how to intervene and support the task can be regarded as a simple and obvious choice; *Uncertainty/Doubt* - The extent to which you have confidence in the machines' decision on when and how to intervene and support the task; *Judgement/Awareness* - The extent to which the machines decision on when and how to intervene and support the task requires assessment, knowledge, and understanding of the task; *Faith* - To what extent you believe that the machine will be able to intervene and support the tasks in other systems states in the future; *Demand for Trust* - Level of trust required from you when the machine intervenes and supports the task; *Supply of Trust* - Level of trust actually provided by you when the machine intervenes and supports task; *SART Demands on Attentional Resources (D)* - How demanding is the task on your attentional resources ? Is it excessively demanding (high) or minimally demanding (low) ?; *SART Supply of Attentional Resources (S)* - How much of your attentional resources are you supplying to the situation? Are you making the greatest, possible effort (high) or giving very little attention (low) ?; *SART Understanding of the Situation (U)*- How well do you understand the situation? Do you understand almost everything (high) or virtually nothing (low) ? (N.B. *SA(c)*, is calculated from the SART ratings using the formula SA(c)= U - (D-S)).

At the end of the experimental session, subjects completed ratings of the system on 21 dimensions of human-computer teamwork, based on the model proposed by Taylor and Selcon (1).

## 3.3 Results

a. *Analysis of Variance.* ANOVAs on the dependent variables showed no clear pattern of effects arising from the manipulation of computer co-operation. Significant differences were found between the scenarios on Resource Management tank level 1 and 2 deviations ($p<0.001$). Newman Keuls tests showed significantly more deviations in the Co-op2 scenario than in the other three scenarios ($p<0.01$). There was a small but significant difference between the subject groups on ratings of the supply of trust ($p<0.05$). There were no significant subject group/scenario condition interactions, and thus no proof of order or transfer effects between the scenario conditions.

b. *Correlations.* Correlation analysis was performed on the performance and ratings data. This analysis revealed significant correlations between the many of the variables. A schematic representation of the significant correlations is provided in Figure 3, following the style used by Riley (6,7). In Figure 3, variables with significant correlations ($r > 0.40$) are linked by lines, with the strength of association indicated by the line width.

c. *Factor Analysis.* Factor analysis of the subjective ratings found that four factors accounted for 62% of the total variance in the data. The results are summarised in Table 1, with ratings variables that obtained significant loadings on the four factors ($> 0.45$) shown in order of reducing weight, with positive or negative values (+/-ve).

| Factor 1 (-ve) 21.55 % Variance. | Factor 2 (-ve) 17.65 % Variance. | Factor 3 (+ve) 11.97 % Var. | Factor 4 (-ve) 10.95 % Var. |
|---|---|---|---|
| Auto. Reliability | Self Accuracy | Impact (-ve) | Uncertainty/ Doubt |
| Auto. Confidence | Confidence | Supply of Trust | Faith |
| Auto. Dependability | SA(c) | Supply of Attent'l | Decision Complexity |
| Auto. Accuracy | Self Confidence | Resources | Demands on Attent'l |
| Auto. Predictability | Accuracy | Demand for Trust | Resources |
| Supply of Attent'l | Understanding of | | |
| Resources | Situation | | |
| Supply of Trust (+ve) | | | |

**Table 1. Factor Analysis of Ratings**

d. *Teamwork Ratings.* Factor analysis of the ratings indicated that the eigen values were too small to warrant further consideration. Mean ratings on the teamwork model dimensions are illustrated in Figure 4, relative to the benchmarks for good and poor teamwork reported by Taylor and Selcon (1). This shows that the ratings fall between the values obtained previously, with the exception of *heterogeneity* of Resources, and *cohesiveness* of Structure, which score below their respective poor teamwork benchmarks. In general, the ratings for the Goals dimensions are relatively high, indicating good teamwork, whereas the ratings for Structure and Processes dimensions are relatively low, indicating the converse.

| | LOW | | HIGH |
|---|---|---|---|
| **GOALS:** Clarity - defined performance objectives | ♦ | ▼ | ♦ |
| Common structure - shared understanding | ♦ | ▼ | ♦ |
| Tracking - awareness of changing objectives | ♦ | ▼ | ♦ |
| Impact - critical for mission success | ♦ | ▼ | ♦ |
| Achievement - high probability of success | ♦ | ▼ | ♦ |
| **RESOURCES:** Sufficiency - enough expertise & ability | ♦ | ▼ | ♦ |
| Availability - readiness for application | | ♦ ▼ | ♦ |
| Heterogeneity - variable/unique expertise | ▼♦ | | ♦ |
| Compatibility - ability to combine & integrate | ♦ ▼ | | ♦ |
| Enhancement - ability to add expertise | ♦ | ▼ | ♦ |
| **STRUCTURE:** Goal driven - governed by performance | ♦ ▼ | | ♦ |
| Accessibility - facilitates access to resources | ♦ | ▼ | ♦ |
| Cohesiveness - attracts conformity to norms | ▼♦ | | ♦ |
| DPA - real-time role & task distribution | ♦ ▼ | | ♦ |
| LOA - degrees of independent functioning | ♦ | ▼ | ♦ |
| **PROCESSES:** Wide bandwidth - many comm. modes | ♦ ▼ | ♦ | |
| Bidirectionality - two way information flow | ♦▼ | | ♦ |
| Shared initiative - leadership turn-taking | ♦ ▼ | ♦ | |
| Common knowledge-shared understanding | ♦ ▼ | | ♦ |
| Trust - willing to accept others' judgements | ♦ | ▼ | ♦ |

Key :▼ = MAT ratings ; ♦  ♦ = Benchmarks

**Figure 4. Teamwork Ratings**

## 4. Discussion

The MAT data indicate that the subjects successfully compensated for the variability in the computer aiding performance, but with little effect on their assessments and attitudes to the automation. The high error rates on the resource management task in the Co-op 2 scenario, where the task remained fully automated, indicates that manual control was capable of producing better performance. Indeed, some subjects reported that the automation performance seemed relatively poor compared to manual. This raises questions about the validity of the MAT automation software. In contrast, despite clear instructions to monitor the imperfect aiding invocation, with regard to timeliness and the level of aiding provided, most subjects seemed unaware of the

Figure 3. Schematic Structure of the Correlated Variables

experimental manipulation of computer co-operation. Only 3 subjects noticed the computer warning to go fully automated, and then failing to do so, which raises questions about the adequacy of the MAT screen format. Mean ratings tended to be higher on four out of five automation performance dimensions (automation accuracy, dependability, reliability, predictability) following the co-operation scenarios, but none significantly so. Trust ratings varied about the middle of the rating scale (mean = 4.23; SD = 1.15). Mean trust supply ratings were higher, and SA(c) means were lower, following the co-operation scenarios, but again the differences did not achieve statistical significance.

Despite the apparent insensitivity of the dependent variables to the experimental conditions, the correlation and factor analysis present a reasonably coherent picture of the relationships between the variables through the course of the experiment. The co-operation manipulation probably exerted little influence over the picture that emerges. Considering first the correlations, the supply of trust was related to confidence in automation performance and to its perceived accuracy, reliability, and dependability. An associated automation performance factor in the factor analysis (Factor 1) accounted for the largest proportion (21.55%) of the variance in the ratings data. The correlations show trust was inversely associated with impact/survivability, or the negative consequences of relying on computer to support the task, i.e. the more adverse the consequences were perceived to be, the less trust was supplied. A similar trust related factor (Factor 3) accounted for 11.97% of the ratings variance. Faith, referring to future performance, was more associated with the requirement for judgement and awareness in the computers decisions, with weak associations with decision uncertainty and doubt, and with perceived automation reliability. A similar, relatively weak uncertainty/faith related factor (Factor 4) accounted for 10.95% of the ratings variance. Self-confidence was linked to assessments of the accuracy of own performance, and to awareness and understanding of the situation. Self confidence/accuracy was associated with a strong factor (Factor 2), which accounted for 17.65% of the ratings data. Understanding of the situation was associated with the MAT performance data. Thus, consistent with the evidence of successful

compensation, performance was associated with awareness and understanding of the task, rather than with the apparent performance of the automation.

Several points follow from the above. It seems likely that subjects experienced difficulty in distinguishing between tasks demands arising from the external situation, and demands associated with variability in the performance of the adaptive aiding. The subjective assessments were surprisingly uncritical of the unco-operation scenarios. There was no substantial loss or gain in trust, and consequently no basis for assessing trust recovery following unco-operative scenarios. Equally, there was no evidence of misplaced trust or of attitude changes to resolve cognitive dissonance. But then initial levels of trust would need to be high for awareness of an inappropriate reallocation to cause a revision of attitudes. It seems likely that the appropriateness of aiding invocation, level of assistance, and reallocation strategy are difficult to judge in a dynamic situation. The demands associated with the automation performance were communicated through a common interface, and shared a common representation with the external task demands. Variations due to the automation could be distinguished only by inference from differences between expected and actual automation performance. It seems likely that the expectations for automation performance were not clearly formulated, and that the compensation occurred smoothly and naturally, without a reluctant hand-over. Subjects appear to have maintained throughout a reasonably high level of confidence in their own ability, and in the computer's assistance, to perform the task. While self confidence is maintained, and compensation occurs without awareness, deviations are likely to be attributed to own, rather than computer, performance. Confidence in adaptive aiding seems dissociated from adaptation performance. Sen and Boe (14) report a similar dissociation between confidence and accuracy in computer-aided decision making. The teamwork benchmarks and ratings indicate that the MAT provides a reasonably representative environment for investigating human-computer teamwork. However, the profile falls considerably short of providing many of the qualities of good human-human teamwork. A more demanding and risky task, performed over a longer duration, with more complex scenarios than used in the

91

present study, may be needed for a more complete picture of what happens to trust following adaptation failure.

## 5. CONCLUSIONS

The First Law of Adaptive Aiding, that computers should be able to take tasks, but not give them, seems sound. Awareness of the current task allocation strategy is an important factor for system effectiveness, but this may not easily be achieved with adaptive aiding. Awareness is needed to avoid task contention, and to ensure that tasks are not overlooked or performed incorrectly. Trust is built on awareness of proven performance. Bi-directional compensation without awareness might engender false trust by others outside the team.

Conflicting or unco-operative intelligence could arise from poor design, such as inappropriate adaptive logic. For this reason, adaptive strategies, such as manual reallocation, will need careful adaptive logic to ensure their appropriateness. Alternative forms of adaptive logic have been proposed based on either critical events, pilot performance measurement, pilot physiological assessment, and pilot modelling. The validity of the threshold criteria for triggering allocation and re-allocation will be critical. Predictions of pilot workload would seem to be the logical candidate for a model-based adaptive logic, particularly for manual task re-allocation. However, sufficiently reliable individual workload threshold criteria will be difficult to obtain from currently available generic workload models and measurement tools. Given the low predictive validity of human performance and workload models, the possibility of an inappropriate re-allocation from an operator model-based logic, or any other logic, will need to be anticipated in the design of adaptive systems. Measures will need to be taken to guard against the consequences of inappropriate allocation, adaptation breakdown or failure, or of what might otherwise appear to be unco-operative automation.

As a safeguard, the system will need to establish operator willingness and readiness to accept tasks before reallocation. In addition, safeguards will need to be built into the human-computer interface to ensure that the operator has the necessary awareness and control of the current functional and task configuration. It may not be sufficient to provide legends for automation status; pictorial representations or dialogue may be needed for comparing the pilot's expectations with EC's plans and intentions. Poor co-operation could lead to mistrust, or be perceived as systematic and intentional, engendering a sense of paranoia. Safeguards may be needed to ensure that the H-EC share a common understanding of all the system meta- and sub-goals underlying co-operative functioning, i.e. the prime directives. It may be time to start considering how to implement the sentiment, if not the substance, of Asimov's Law's of Robotics(15)[1].

## 6. REFERENCES

1. Taylor, R.M. and Selcon S.J., "Operator and Automation Capability Analysis: Picking the Right Team". AGARD CP 520, April 1993, Paper 20.

2. Muir, B.M. "Trust Between Humans and Machines, and the Design of Decision Aids. Int. J. Man-Machine Studies, 27, 1987, pp 527-539.

3. Lerch, F. and Prietula, M. "How Do We Trust Machine Advice?" In "Designing and Using Human-Computer Interfaces and Knowledge-Based Systems". Salvendy, G. and Smith, M. Eds. 1989, Amsterdam, Elsevier..

4. Parasuraman, R., Molloy, R., and Singh, I. "Performance Consequences of Automation-Induced Complacency". Int. J. Av. Psych., 3, 1, 1993, pp 1-23.

5. Taylor, R.M. "Trust and Awareness in Human-Electronic Crew Teamwork". In " The Human-Electonic Crew: Can They Work Together?" WRDC-TR-89-7008, Wright Patterson AFB, OH., 1988.

6. Riley, V. "Modelling the Dynamics of Pilot Interaction with an Electronic Crew". In " The Human-Electonic Crew: Is the Team Maturing?" WL-TR-92-3078, Wright Patterson AFB, OH, July 1992, pp 103-107.

7. Riley, V. "A Theory of Operator Reliance On Automation". In "Human Performance in Automated Systems: Current Research and Trends", Mouloua, M. and Parasuraman, R. (Eds), Hillsdale, New Jersey, Lawrence Erlbaum Associates, 1994.

8. Lee, J.D. and Moray, N. "Trust, Control Strategies and Allocation of Function in Human Machine Systems". Ergonomics, 35, pp 1243-1270.

9. Rouse, W.B. "Twenty Years of Adaptive Aiding: Origins of the Concept and Lessons Learned". In "Human Performance in Automated Systems: Current Research and Trends", Mouloua, M. and Parasuraman, R. (Eds), Hillsdale, New Jersey, Lawrence Erlbaum Associates, 1994, pp 28-32.

10. Carmody, M.A. and Gluckman, J.P. "Task Specific Effects of Automation and Automation Failure on Performance, Workload and Situational Awareness", Proc.of 7th Int'l Symp. on Av. Psych., Ohio State Univ., 1993, Vol 1, pp 167-171.

11. Parasuraman, R., Mouloua, M., Molloy, R., and Hilburn, B. "Adaptive Function Allocation Reduces Costs of Static Automation", In "Proc. of 7th Intern'l Symp. on Av. Psych., Ohio State Univ., 1993, Vol 1, pp 178-181.

12. Comstock, R.J. and Arnegard, R.J. "Multi-Attribute Task Battery. Preliminary Technical Report", NASA Langley Research Centre, Hampton, VA.

13. Taylor, R.M. "Situation Awareness Rating Technique (SART): The Development of a Tool for Aircrew Systems Design". AGARD CP 478, April 1990, Paper 3.

14. Sen, T. and Boe, W.J. "Confidence and Accuracy in Judgements Using Computer Displayed Information". Beh. and Info. Tech., 10, 1, 1991, pp 53-64.

15. Asimov, I., "I, Robot", Grafton, London, 1967.

---

[1] 1. A robot may not injure a human being, or, through inaction allow a human being to come to harm.
2. A robot must obey the orders given it by human beings except where such orders would conflict with the First Law.
3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.

# Communication in the Human-Electronic Crew: What can we learn from human teams?

A.T. Lucas,        Logica UK Ltd (Cambridge Division)
S.J. Selcon [1],   DRA Farnborough (MMI Dept)
A.W. Coxell,       DRA Farnborough (MMI Dept)
H.J. Dudfield,     DRA Farnborough (MMI Dept)
N. O'Clarey,       Logica UK Ltd (Cambridge Division)

## Abstract

This paper reports findings from studies addressing the communication requirements of aircrew in current operational two-seat aircraft, in an attempt to identify issues relating to future single-seat operations. A communications analysis was conducted on video-taped air-to-ground training missions flown by experienced aircrew in a Tornado flight simulator. The analysis categorised the communications as Statements, Assertions, Questions, Confirmations, and Discussions, allowing a quantitative assessment of the data; it was supported, where appropriate, by findings from an information requirements analysis conducted through semi-structured interviews with RAF aircrew.

The results are evaluated within the teamwork model of Taylor and Selcon (1993), who identified communication processes as the primary area of concern in the development of human-electronic teams, and their implications are discussed for issues of bi-directionality of communication, shared initiative, common knowledge, and trust. The data showed that more communication was initiated by the navigator, with the content of that communication comprising mainly statements to maintain current and common knowledge and situational awareness. They also showed a high degree of two-way information flow and initiative-taking. Further, the absence of contradictions, and the relatively small number of questions and discussions, implied that a high degree of trust existed between the crew. Recommendations for designing systems to support successful communications between human and electronic crewmembers are drawn.

[1] Requests for reprints should be sent to Dr Steve Selcon, Human Factors Group, MMI Dept., Aircraft Sector, Defence Research Agency, Farnborough, Hants GU14 6TD.

## Introduction

The trend towards single-seat crewing for future military aircraft has required consideration to be given to the use of an Electronic Crewmember (EC) to replace the functions of a second human crewmember. The concept of man and machine working as an intelligent, co-operative team is considered by many as being central to the production of an effective EC (1; 2). One potential solution to designing systems that will facilitate human-machine teamwork is to use expert or Knowledge Based Systems (KBS). The implementation of expert systems in such a way as to produce a synergistic relationship with the human operator raises a number of Human-Computer Interface (HCI) issues, the resolution of which is likely to be crucial in any successful system. Communication via the HCI is considered to be paramount in the design of interactive computer systems. The use of KBS technology in the cockpit allows

the system, or the designer through the system, to use knowledge of the user's information requirements to structure the interface to meet the needs of the operator most adequately. One method of achieving this is to employ an expert system containing knowledge about goals, actions available, etc. This method would support the interaction of the operator with both his expert and non-expert systems, by providing dynamic sensor-fused and information-fused displays. The embedded knowledge the expert system contains allows the display of information relevant to the task being performed, thus providing the information when, and in the form, required. An example of such an interface is the Pilot Vehicle Interface (PVI) described by Shellnut (3).

To identify in more detail the communication needed for effective teamwork, a model of teamwork is required. Such a model was proposed by Selcon & Taylor (4) and is shown in Figure 1. An attempt to validate the model was conducted by Taylor & Selcon (5), who used a teamwork audit technique requiring subjective estimates on 20 teamwork dimensions, to identify which areas of teamwork were relevant and important in successful teaming for both two-seat human-human (pilot-navigator) and single-seat human-machine (pilot-avionics systems) crews. Although 15 of the audit constructs were sensitive to differences in realistic mission teamwork scenarios across the two types of team, only four showed significant differences between the single- and two-seat crews. The four constructs, which all lay within the Processes stage of the Selcon & Taylor (4) model, were: Bidirectionality of communication; Shared Initiative; Common Knowledge; and Trust. Since these constructs indicate the shortfall in current human-machine teamwork, and are all dependent to some extent at least on communication, the results of the present study will be used to attempt to identify possible solutions to such shortfalls and to suggest how such solutions could be made relevant to designing an effective human-electronic team. The approach taken here is to examine how communication is carried out in effective human teams (in this case military pilot-navigator teams) and to attempt to draw inferences from the findings for the communication requirements which would be needed to allow successful human-electronic teamwork to take place.

Figure 1 - Teamwork Model



94

## Method

Ten RAF Tornado aircrew (five pilots and five navigators) were used in the study. Data were collected by the post-hoc analysis of video-taped missions flown in a realistic Tornado simulator. Subjects were not aware,when performing the missions, that analysis of their communications would be conducted. Each mission consisted of two separate ground attack scenarios performed by a pilot-navigator team.. Enroute threats appeared on course to targets. Each mission was terminated by the introduction of an emergency situation. Thus, although the missions could not be made identical (since each team would approach the task slightly differently), all contained equivalent tasks/task difficulty.

To quantify the communication between crew members during the mission, each vocal output was categorised under one of five headings:

i) Statement - an unprompted report of current situation, status, or data.

ii) Confirmation - a response to a request for data/information.

iii) Assertion - an unprompted declaration that an action is being or should be taken.

iv) Question - a request for information.

v) Discussion - a verbal interchange to examine or debate a problem.

Communication was analysed by means of a cumulative total of the number of each type of output, for each crewmember, through the mission.

## Results

The total numbers of each type of output, split by crewmember, can be seen in Figure 2. Statistical analysis was conducted on the scores using a two-way Analysis of Variance (ANOVA) with a between-subjects factor of crew type (two levels: Pilot and Navigator) and a within-subjects factor of communication type (five levels: Statements, Confirmations, Assertions, Discussions, Questions). Post-hoc analysis carried out using t tests, adjusted using the Bonneferroni inequality to produce an experimentwise error of less than 5%.

The main effect of crew type approached significance ($F_{1, 8} = 5.3$, $p = 0.0505$), with navigators (mean = 11.9) producing more vocal outputs than pilots (mean = 5.8). The failure quite to reach significance at the 5% level is likely to be a result of the low power of the test available from such a small sample size (n = 10). There was a highly significant main effect of communication type ($F_{4, 32} = 16.6$, $p < 0.001$). Post-hoc tests showed that the difference was produced by the higher mean for Statements ($p < 0.001$) than for the other four types of communication. There were no reliable differences between Confirmations, Assertions, Discussions, and Questions, and the interaction between crew type and communication type was not significant ($F_{4, 32} = 2.1$, $p > 0.05$).

Figure 2 - Total number of each type of vocal output for pilots and navigators

## Discussion

It can be seen from the above results that significantly more statements are made than any other form of vocal output. Further, the absence of an interaction effect implies that this is true of each crew member. Figure 2 shows that statements account for more of the outputs than all the other categories combined. This finding has implications for the design of an EC. Since statements are merely reports of factual occurrences, they are amenable to generation by an EC. However, they are context dependent, and reflect intent inferencing on the part of the speaker.

The results also imply that a much greater amount of communication is initiated by the navigator than the pilot. This in turn implies that, if the EC is to take over the functional responsibilities of the navigator,

great emphasis must be placed on the communication abilities of the EC, particularly since the majority of communication is unprompted. Thus for an EC to become a successful team member, it must be able to initiate appropriate communication on a regular basis. Care must be taken, however, in drawing such conclusions, since the main effect of crew type narrowly failed to reach significance in this study. Also, it should be borne in mind that this finding may reflect the way information is displayed in the Tornado, rather than a generalisable trend across all two-seat aircraft. If certain information is displayed only to the pilot or the navigator, then communication is being forced to occur. If all information were available to both team members, then a different communication pattern might be observed.

The results of this study can also be used to examine the findings of Taylor & Selcon (1993) on teamwork differences between human-human and human-machine crews. Each of the four dimensions of teamwork found to be sensitive in the Taylor & Selcon study will be evaluated separately. The first of these dimensions was bidirectionality of communication. In other words, for a team to be effective, communication must be a two-way process. Current human-machine teams failed to show such bidirectionality. The results of the present study support the requirement for information flow from both team members. The high number of vocal outputs from the navigator underlines this, since, if the EC is to replace primarily the navigator's functions, it must be able to produce such regular communication to support the pilot's current mental model of the situation and hence his situational awareness.

A second dimension where current human-machine teams showed weaknesses was in shared initiative. The results of this study show that, since the majority of communication was both unprompted and instigated by the navigator, the EC must be able to take the initiative in providing information when and where appropriate. As mentioned previously, this is likely to require the ability to assess the actions being taken by the pilot, and to infer from them his current intent, to maintain the relevance of the EC-generated communication. Such intent inferencing will be dependent on the third dimension discussed by Taylor & Selcon, namely common knowledge. If common knowledge is not designed into an EC, the

implementation of current human-human communication patterns is likely to be difficult to achieve. Navigators have, through training and operational experience, a good understanding of the mission tasks being performed by the pilot. Such understanding is critical in allowing autonomously generated communication to be both effective and minimal. Thus, for the EC to achieve the type of communication patterns found in this study, knowledge of what the human needs to know at a given stage of the mission will be needed.

The fourth dimension on which single- and two-seat crews differed was trust. Although trust per se is related only tangentially to communication, the results of this study do have some relevance to the consideration of human-EC trust. The absence of negations and the very limited numbers of discussions within the human team members in this study imply that a high level of trust must exist (at least in terms of verbal communication). Without such trust, communication will be much less efficient since checking/querying of information being conveyed will be required. Similarly, without effective and regular communication, appropriate levels of trust will be less likely to develop since each team member will have less understanding of what the other team member knows and is doing. Thus trust and communication must be considered together in the design of a successful EC.

**Conclusions**

The above discussion has identified a number of lessons that can be learnt from human-human communication, relevant to the

design of an effective EC. For example, communication must be available in both directions, but with emphasis on that provided by the EC. The EC must be able to take the initiative in providing information without being asked for it. Further, there will be a requirement for intent inferencing in the EC, through common knowledge with the human, to maintain the appropriateness of such communication. Finally, the design of an EC that is able to communicate effectively will have benefits in producing trust between human and electronic team members. The results of the present study also indicate the potential utility of the approach of examining human communication as a means of identifying communication requirements for human-electronic teamwork. Future areas where such an approach could possibly be used would include analysis of the common knowledge required to allow intent inferencing, and investigation of other factors affecting trust between human-electronic crews.

## References

1. Emerson, T. J., Reinecke, M., Reising, J. M., & Taylor, R. M. (Eds.). (1988). The Human-Electronic Crew: Can They Work Together? Proceedings of a joint GAF/USAF/RAF Workshop. Farnborough, UK: RAF Institute of Aviation Medicine, BSD-DR-G4.

2. Emerson, T. J., Reinecke, M., Reising, J. M., & Taylor, R. M. (Ed.). (1990). The Human-Electronic Crew: Is the Team Maturing? Proceedings of the 2nd joint GAF/USAF/RAF Workshop. Farnborough, UK: RAF Institute of Aviation Medicine, PD-DR-P5.

3. Shellnut, J. B. (1988). Pilot Vehicle Interface Management. In T. J. Emerson, M. Reinecke, J. M. Reising, & R. M. Taylor (Eds.), The Human-Electronic Crew: Can They Work Together? Proceedings of a joint GAF/USAF/RAF Workshop, BSD-DR-G4. Farnborough, UK: RAF Institute of Aviation Medicine.

4. Selcon, S.J. & Taylor, R.M. (1991). Psychological principles for Human-Electronic crew teamwork. In Emerson, T. J., Reinecke, M., Reising, J. M., & Taylor, R. M. (Ed.) The Human-Electronic Crew: Is the Team Maturing? Proceedings of the 2nd joint GAF/USAF/RAF Workshop. Farnborough, UK: RAF Institute of Aviation Medicine, PD-DR-P5.

5. Taylor, R.M. & Selcon, S.J. (1993). Operator and automation capability analysis: Picking the right team. In Proceedings of the NATO AGARD Joint FMP/GCP Symposium on Combat Automation for Airborne Weapons Systems, CP 520, 20. Neuilly-sur Seine: NATO AGARD.

# Improving Communication and Trust with Memory Techniques

**D.A. Bekerian and J.L. Dennett**
MRC Applied Psychology Unit
Cambridge, England

## 1. SUMMARY

Communication is one of the most essential ingredients for successful team work. Sharing knowledge, that is communicating to other team members what you know, is critical. When communication involves remembering a specific event, or details that happened during the execution of some procedure, psychological data on memory-enhancing techniques can be of great use. The discussion presented here considers one very successful memory technique, known as the Cognitive Interview technique(CI). CI has already proven itself, in the United States, Continental Europe and Great Britain, as a powerful communication aid, in contexts ranging from evidential interviews with trauma victims to marketing research on people's shopping habits. CI, it would seem, improves communication. We provide some background information about CI, its basic memory techniques, and also discuss how other features of communication, like trust or rapport building, can be improved through the application of memory techniques like CI.

## 2. INTRODUCTION

Failures in communicating information -- particularly in the context of team operations -- lead to difficulties. Consequently, any technique that can enhance the successful communication of information between team members may become very important as communication tools. In some situations (for example, debriefing procedures) a person may be required to communicate details about specific events, or specific procedures executed, for example personal experiences regarding the execution of some operations. Although many of these details might be recorded, through technological procedures , the technological evidence may be absent or faulty, so accurate recall becomes essential. Under these circumstances,, cognitive psychology can function as a source of information for techniques that can be used to enhance memory, and thereby enhance communication. Today, we will be providing you with information concerning one particularly successful approach, that called the Cognitive Interview technique (Bekerian & Dennett, 1992).

## 3. THE COGNITIVE INTERVIEW TECHNIQUE

The Cognitive Interview technique (which we will refer to as CI) has been one of the most effective applications of memory theory to real world problems, particularly when a person is attempting to remember details from a specific event.

CI is used widely by police forces both here in the States and also in Great Britain and Europe. The technique, originally devised by Ed Geiselman (UCLA) and Ron Fisher (Florida), has been extensively researched, across different paradigms, materials, retention intervals and subject populations -- to name only a few of the variables that have been studied. The general finding, although there are exceptions, is that CI increases the amount of correct recall without simultaneously increasing the number of errors. CI is distinguished from other interview techniques by its emphasis on four cognitive mnemonics. We present a brief description of each.

## 4. THE COGNITIVE MNEMONICS OF CI

### 4.1 Reinstate the context.
*The individual is asked to try and reinstate (or recreate) as much of the surrounding context that occurred during the event, including emotional/internal states.*

Generally, context reinstatement gives the person to additional cues that will enhance their ability to remember details. This is a phenomenon that has had extensive experimental investigation in the psychological literature on memory (e.g., Tulving and Thomson, 1973; Cutler, Penrod & Martens, 1987; Malpass & Devine, 1981; Smith, 1979; see also Davies & Thomson, 1988). Generally, the effects are particularly notable when someone is being asked to produce a narrative report, such as when someone is asked to report an incident, or say what happened, including details of people, actions and objects.

There are any number of ways that features of the original context might be reinstated. For example, a person can be asked to image, mentally, the surroundings before recounting an event; they can be asked to recreate any feelings or reactions that they had at the time; they can be asked to form images of objects, or of actions done, before they attempt to report details; they can be asked to remember the order of a sequence of actions (e.g., Eldridge, Barnard & Bekerian, 1993); or they can be taken back to the scene of an incident (e.g., Bowers & Bekerian, 1984). As of yet, there has been no systematic study that has determined which type of technique is more effective overall, or when specific technique might be more effective than others. However, as we have already mentioned, the effects are particularly noticeable when someone is being asked to produce a narrative report.

### 4.2 Report Everything.
*The individual is encouraged to report as much as s/he can remember, irrespective of the level of subjective confidence associated with the information. or the perceived importance of the information.*

Many people believe that confidence is a reliable indicator of the accuracy of information (see Noon & Hollin, 1987). As a consequence, people may edit what they report because they feel uncertain about a piece of information (see Loftus, 1978). However, the empirical work suggests that the relationship between confidence and accuracy is unreliable (e.g., Kassin, Ellsworth & Smith, 1989, Wells & Lindsay, 1985). Consequently, it is important to try and force individuals to disregard the confidence with which they recall a piece of information, and to report everything even if they feel uncertain. Oftentimes, they might be correct, but not terribly confident.

Additionally, many people will fail to report a detail because they think that it is irrelevant or unimportant. Of course, it is often the case that seemingly irrelevant details can, in fact, be extremely useful and important. Consequently, it is critical that the person be encouraged to report all details, irrespective of the perceived importance of the information.

The theoretical rationale for the use of this strategy is compatible with the work on Signal Detection Theory, which was originally devised for modelling signal/noise detection. (Signal Detection Theory suggests that a person's willingness to report an event will depend upon the particular response bias (or confidence level) that is adopted. )

### 4.3 Recall Events in Different Order.
*The individual is asked to recall the event in more than one sequential order, that is, recall the event from the end to the beginning, or from the middle, etc. This strategy can also be applied to lower level details, as in the scanning of images.*

According to memory theory, a change in the order of recall results in the person remembering the event from different mental perspectives; as these perspectives change, the nature of the information available for the report also changes, resulting in the recalling of new information. Empirical evidence has shown that changes in recall order can influence memory for information from different sequential positions, for example memory of actions occurring during the middle of an event is enhanced if backward serial probing is used (e.g., Geiselman & Callot, 1990; Mingay, Dennett & Bekerian, 1984).

### 4.4 Repeated Recall Attempts.
*The person is asked to recall the event more than once, with some short delays between recall attempts. . The person is encouraged to disregard what they previously remembered, and try and remember the event itself*

The assumption is that repeated recall will enable the person to adopt if you will different mental perspectives, thereby increasing the chances that new information will be reported   Empirical and field studies support this assumption (e.g., Shaw and Bekerian, 1992). It is essential,

though, that the person is encouraged to remember the event itself, rather that what was previously said, otherwise the person simply attempts to be consistent with previous recall attempts and nothing is gained.

### 5. RELATIVE EFFECTS OF THE DIFFERENT MNEMONICS
There have been investigations that have attempted to examine the relative effectiveness of the four CI mnemonics. Geiselman et al. (1986a) reasoned that if individual strategies were as effective as the complete CI procedure, CI could be modified or shortened, thereby saving considerable time for professional investigators. They found that context reinstatement or instructions to be complete were more effective than standard interview instructions; but, that neither were as effective as the combined CI procedure (see also Bekerian, Dennett, Hill & Hitchcock, 1990). Similarly, George (1991) has discovered that context reinstatement is the easiest of all mnemonics for the interviewee to use consistently and, that it is the easiest mnemonic to administer.

### 6. COMMUNICATIVE SKILLS OF THE INTERVIEWER
In addition to the emphasis on cognitive mnemonics, CI equally emphasises the quality of communication between the interviewer and the interviewee, e.g., the importance of rapport building. Fisher and Geiselman (1992) also state that it is important for both the interviewer and the interviewee to share in their expectations of why and how the interview is being conducted. These suggestions appear in most guides to interviewing and, consequently, are not unique to CI (see Shepard, 1988). Nonetheless, they are critical suggestions and should not be underestimated: The fact that both participants share a common goal, have similar expectations, and are seen to be co-operating is one of the most critical features of good communication, and will directly influence the accuracy and quality of any report or account that is produced.

Finally, other guidelines about interviewing that are advised by CI proponents, although they are not considered to be integral to the administration of CI techniques (see Fisher & Geiselman, 1992), include the preference of open-ended questions over closed questions; the care needed in asking questions at a rate which allows the person to respond; the need to use language that is compatible with that used by the person (in other words, do not re-describe something using your own terms).

### 7. AREAS THAT MAY REQUIRE SOME TROUBLE SHOOTING
Fisher et al. have also provided some useful information about problems that seem to arise with alarming regularities across formal interviews, in their case police interviews. We wish to stress to you that these problems are not restricted to police interviews, and quite regularly occur in most interviews, be they job interviews, interviews with one's doctor or interviews with witnesses to crimes.

We stress that these problems -- which reflect common interview practices -- are noted in the most experienced of interviewers. So, experience does not necessarily make an interviewers immune to these problems.

Three particularly common interview practices were noted as likely to actually inhibit witnesses from producing a complete and coherent account. They are a) excessive interruptions; b) excessive use of a question-answering format; and, c) inappropriate sequencing of questions. I go through these very briefly in turn. First of all, excessive interruptions.

The act of remembering requires some concentration. Any interruptions will, by definition, distract the witness. Most experienced interviewers begin an interview by asking for a free narrative account. However, Fisher et al. (1989), noted that interviewers on average interrupted four times per every response given by the witness (an average of every 7.5 secs.). George (1991) also notes in his field study of English Constabularies that officers not trained in CI asked significantly more questions per 10 minutes of interview than officers trained in CI.

The second common problem noted by Fisher et al. is the excessive use of specific question-answering formats. This is asking questions where the answer is either yes/no, or the choice of answers is severely restricted, such as did he do X or Y? Certainly, all interviews rely on a question-answering format. However, Fisher et al. comment that the format often is one where very specific questions are posed, rather than more general ones. Although there are benefits to short-answer questions, they may inhibit the person from remembering additional information.

The third most common problem is the inappropriate sequencing of questions. This takes a number of forms. For example, interviewers sometimes follow a standardised, rather formalised, line of questioning, rather than tailoring the questioning for the specific person in question. The result is that the line of questioning appears to be motivated by a specified checklist, rather than motivated by the account produced by the person.

Another example of inappropriate sequencing is when interviewers ask questions that are pertinent to information already given by the person, rather than pertinent to information currently being given, e.g., returning to a previous point while the person is trying to answer another question. This interrupts the person's train of thought. Certainly an interviewer may wish to clarify on a point mentioned earlier. But, such interruptions are often not timed sensitively and the interviewers merely interrupts the person.

Finally, interviewers may be very unstructured in their line of questioning. One question may ask about a perceptual features of an event, then the following question will ask about an entirely different type of feature, such as the order in which an action took place. This rapid and unstructured shift in topic prevents the person from concentrating his/her attention on producing an accurate and comprehensive account.

## 8. RELATIVE EFFECTS OF THE DIFFERENT MNEMONICS

There have been investigations that have attempted to examine the relative effectiveness of the four CI mnemonics. Geiselman et al. (1986a) reasoned that if individual strategies were as effective as the complete CI procedure, CI could be modified or shortened, thereby saving considerable time for professional investigators. They found that context reinstatement or instructions to be complete were more effective than standard interview instructions; but, that neither were as effective as the combined CI procedure (see also Bekerian, Dennett, Hill & Hitchcock, 1990). Similarly, George (1991) has discovered that context reinstatement is the easiest of all mnemonics for the interviewee to use consistently and, that it is the easiest mnemonic to administer.

## 9. RESTRICTIONS ON THE USE OF CI

CI is really only effective when some specific event, or events, are being remembered. Thus, for remembering facts, like telephone numbers, CI would be completely inappropriate. CI may also be more effective when the event is relatively infrequently experienced. Thus, should the event in question be one that the person has experienced on a regular and consistent basis (such as returning home from work each day), it is important that certain steps be taken in order to facilitate the person's use of the mnemonics, for example reinstating as much of the physical environment as possible. Equally, if there had been a delay between the event and the time of recall, the person might be asked to link the event with other significant events that may have happened around that time (a procedure that is called anchoring).

## 10. CONCLUSIONS

CI is one of the most powerful interviewing tools that is available from the psychological literature on memory. CI incorporates basic principles that are known to facilitate the recall of specific events, and also emphasises the importance of communicative skills of the interviewer. In this way, it can be seen as an important communication aid, both in terms of facilitating memory, and also in terms of enhancing trust and rapport between interviewer and interviewee or crew member and debriefing officer. However, it is advisable that prior to the use of CI, interviewers are trained formally. This is primarily because, as is the way with all skills, CI is better administered by people who have had at least some experience in interviewing with CI. Training courses provide such initial experience. Fortunately, there are many courses now on offer internationally that provide such training. So, CI is not only effective but also accessible.

# 11. REFERENCES

Bekerian, D.A. & Dennett, J.L. (1992). The Cognitive Technique: Reviving the Issues. Applied Cognitive Psychology, 7, 275-297.

Bekerian, D. A. & Dennett, J.L. Hill, K. & Hitchcock, R. (1991) Effects of Detailed Imagery on Simulated Witness Recall. In F. Losel, D. Bender & T. Bliesener (Eds.), Psychology and Law Facing the Nineties. 2 vols. Amsterdam: Swets & Zeitlinger.

Bowers, J.M. & Bekerian, D.A. (1984). When will post-event information distort eyewitness testimony? International Journal of Applied Psychology, 69, 466-472.

Cutler, B., Penrod, S. & Martens, T. (1987). Improving the reliability of eyewitness identification: Putting context into context. Journal of Applied Psychology, 72, 629-637.

Davies, G.M. & Thomson, D.M. (Eds.) (1988). Memory in Context: Context in Memory (pp. 339-344). John Wiley & Sons: Chichester.

Eldridge, M., Barnard, P. & Bekerian, D.A. (1993). Schematic influences in

autobiographical memory. MS submitted for publication.

Fisher, R., Geiselman, R.E. & Amador, M. (1989). Field tests of the cognitive interview: Enhancing the recollection of actual victims and witnesses of crime. Journal of Applied Psychology. 74(5) 722-727.

Fisher, R. P. & Geiselman R. E. (1992). Memory-Enhancing Techniques for Investgative Interviewing: The Cognitive Interview. Illinois: Charles C Thomas.

Geiselman, R.E., Fisher, R., MacKinnon, D. & Holland, H. (1986a). Enhancement of eyewitness memory with the cognitive interview. American Journal of Psychology, 99, 385-401.

Geiselman, R.E. & Callot, R. (1990). Reverse and forward recall of script-based texts. Applied Cognitive Psychology, 4, 141-144.

George, R. (1991). A field and experimental evaluation of three methods of interviewing witnesses/victims of crime. Master's thesis, unpublished. University of East London.

Kassin, S, Ellsworth, P. & Smith, V. (1989). The "general acceptance" of psychological research on eyewitness testimony. American Psychologist, 44, 1089-1098.

Malpass, R. & Devine, P. (1981). Guided memory in eyewitness identification. Journal of Applied Psychology, 66, 343-350.

Mingay, D., Dennett, J.L. & Bekerian, D.A. (1984). Memory for a staged incident. Forum, 17, 58-60.

Noon, E. & Hollin, C. (1987). Lay knowledge of eyewitness behaviour: A British Survey. Applied Cognitive Psychology, 1, 143-153.

Shaw, G. & Bekerian, D.A. (1991). Hypermnesia for high imagery words: The effects of interpolated tasks. Memory and Cognition, 19, 87-94.

Shepard, E. (1988). Developing interview skills. In P. Southgate (Ed.). New Directions in Police Training. London: HMSO.F.

Smith, S.M. (1979) Remembering in and out of context. Journal of Experimental Psychology: Human Learning and Memory, 5, 460-471.

Tulving, E. & Thomson, D. (1973). Encoding specificity and retrieval processes in episodic memory. Psychological Review, 80, 352-373.

Wells, G.L. & Lindsay, R.C.L. (1985). Methodological Notes on the Accuracy-Confidence Relation in Eyewitness Identifications. Journal of Applied Psychology, 70(2), 413-419.

# Must the Human - Electronic Crewmember Team Pass the Turing Test?

**Dr John M. Reising**
**Wright-Patterson Air Force Base**
**Ohio, USA  45433**

## 1. SUMMARY

"Somewhat paradoxically, machines that can do more, and do it faster, provide the basis for systems that are increasingly demanding of the human operator, particularly in terms of cognitive requirements" [Howell , 1993, p.235].

The Human - Electronic Crewmember Team is designed to solve this paradox  by having the Electronic Crewmember support the pilots, thereby making their job easier. This support will, in turn, build up the pilots' trust in the Electronic Crewmember. However, this trust may not transfer to others outside the team who realize that one of the members is a machine.   If the machine could appear to behave as a human -- pass the Turing Test -- then trust by outsiders may increase. This paper describes both the Turing Test and guidelines for building trust in human-intelligent machine systems.

## 2. INTRODUCTION

An inherent mistrust of computers has been expressed by many people.   In the development of a teleterminal (a computerized telephone) researchers found that subjects using the prototype said they disliked the software designers they called the "programming priesthood", and by implication, their creations (Hagelbarger & Thompson, 1983). This mistrust could easily extend to the Human - Electronic Crewmember (H-EC) Team, especially the electronic portion of the team.  One could postulate that mistrust might be lessened if the EC seemed more human to those interacting with the H-EC Team. That is, would the team engender more trust if <u>both</u> members appeared human -- if the EC could pass the Turing Test?

In the 1950's Alan Turing wrote an article concerning the intelligence of machines. The issue centered on proving that a machine possessed intelligence, and Turing proposed a test which he hoped would settle this point.  The essence of the test involved an interrogator and two others who were being questioned, none of whom could see each other. The interrogator sat at a computer terminal and sent questions to the other two (one female and one male) who were in separate room and also had computer terminals.  The goal of the interrogator was to determine the gender of each. To make the test more interesting, one of people being questioned was told to tell the truth, while the other was told to give false or misleading answers. Let us say in our case the male was to lie and the female was to tell the truth. For example if the interrogator asked them if they were male, the female would say no, but the male would also say no. Turing further proposed

that if a machine were substituted for one of the humans being questioned, and the interrogator could not tell the difference, the machine would possess the qualities of intelligence. One could speculate that if the EC passed the Turing Test, the team would be trusted more by those who interact with it.

Another major component required for outsiders to trust the H-EC team is that it consistently make correct decisions and execute correct actions (Muir, 1987). The EC's passing the Turing test does not guarantee that the EC will perform flawlessly. Recall that in the Turing Test, after the computer had been substituted for one of the two humans being questioned (say the one who could give false or misleading answers) the object of the interrogator was to determine which of the two was human and which was machine. Turing uses the example of the interrogator asking the two being questioned to add 34957 and 70764. The computer pauses about 30 seconds and presents the answer 105621. Not only did the computer put in an excessively long delay (for a computer), but also came up with an incorrect answer -- 105721 is correct. The reason for these actions is because the computer was trying to mimic human behavior, and, in this case, chose behavior which would definitely not lead to trust by those outside the H-EC Team.

If passing the Turing Test does not build up trust between the team and outsiders, then what are the means to achieve this goal? Trust comes about when an observer sees consistency (and correctness) in behavior from an individual or a team. The team's consistency within itself is built up when the individual team members are themselves consistent and correct in their decisions and actions. The following guidelines discuss how trust can be built within the team, after which the team can then behave consistently to outside observers and gain their trust.

## 3. TRUST GUIDELINES

Guidelines for building up trust within the H-EC Team have been previously presented (Emerson and Reising, 1990; Reising, Emerson, & Munns, 1993). Two of the most important were the establishment of Prime Directives and Levels of Autonomy. Prime Directives are overall governing rules which bound the behavior of the EC so that the pilot does not experience any surprises. Levels of autonomy also bound the behavior of the EC by limiting its decision authority to a level specified by the pilot. These, and other guidelines are discussed in more detail in the following sections.

### 3.1 Define the EC's Prime Directives

One essential feature of a successful team is trust in the other partner. This in turn implies that the partner behaves in a rational and reliable manner; one partner cannot initiate actions which, even though they are logical to it, appear to be illogical to the other. In order to avoid arbitrary

actions, there must be some overall governing rules which provide the logical structure under which both members operate. As examples of explicitly stated governing rules, consider the three laws of robotics (Asimov, 1950).

1. A robot may not injure a human being, or through inaction, allow a human being to come to harm.

2. A robot must obey the orders given to it by human beings except where such orders would conflict with the first law.

3. A robot must protect its own existence as long as such protection does not conflict with the first or second law.

These rules provide the guidance required to allow a robot to perform its job in a reasonable and consistent manner. If the word "pilot" is substituted for the word "human" in the above example, a possible basis for governing the behavior of the EC exists. The three laws stated above are only examples of governing rules, and they would require major changes to be applicable in a military setting. For instance, without modification the ideal robot would not allow pilots to take off, knowing that they were deliberately going in harm's way. The point is, however, that rules of this type provide the basis for consistent behavior by the EC and thereby provide a foundation of trust for the pilot. It is through this trust that an effective team can be built.

## 3.2 Specify the EC's Levels of Autonomy

Another means of establishing pilot trust in the EC is to allow the pilot to decide how much authority or levels of autonomy (LOA) to give the EC. "LOA defines a small set ('levels') of system configurations, each configuration specifying the degree of automation or autonomy (an 'operational relationship') at which each particular subfunction performs. The pilot sets or resets the LOA to a particular level as a consequence of mission planning, anticipated contingencies, or inflight needs" (Krobusek, Boys, & Palko, 1988, p.124). For instance, the pilots could establish a "contract" with the EC in the pre-mission phase. They could, through a dialogue at a computer workstation, define what autonomy they wish the EC to have as a function of flight phase and system function. As an example, weapon consent would always remain exclusively the pilots' task, but reconfiguration of flight control surfaces to get the best flight performance in the event of battle damage would be the exclusive task of the EC.

## 3.3 Conform to the Pilot's Mental Model

The EC must not only be bounded in the overall authority it has, but also must appear to perform logically within those bounds. Mental models play an important part in the efficient operation of systems (Wickens, 1992). Since direct views of the inner workings of a system are often not possible, e.g., the flow of electrons inside the avionics system, displays are a major means of conveying information on the

operation of a system. The closer the EC's display formats conform to the pilots' mental model, the more beneficial they will be. Pilots form a mental picture of how a system should work (at a top level) and base their trust in the system according to how the system conforms to this picture or mental model. A mental model is a representation formed by a user of a system and/or task, based on previous experience as well as current observation, which provides, (most if not all) of their subsequent system understanding and consequently dictates the level of task performance (Wilson & Rutherford, 1989).

These three ideas have been underlined in the above definition to stress its key aspects: representation, understanding and task performance. The pilots' representation leads to their understanding of the system which in turn leads to their performance with the system. For example, if the pilots' mental model of a fuel system pictures the flow valve lever in line with the flow when the fuel is moving and at right angles when the flow is shut off, then that is the way it should be portrayed. It is not important that the valves are electronic and do not have a flow valve handle to turn. An example of not conforming to an operator's mental model is illustrated by the use of reverse notation on early calculators. To add 3 + 2 = 5 instead of punching the keys in this order, the task must be performed in the following order 3 then 2 then +. Needless to say many operators had difficulty in using these calculators.

## 3.4 Make the Interface with the EC Transparent

When pilots communicate with their team members in aircraft such as the Royal Air Force's GR-1 Tornado or the US Air Force's F-15E Strike Eagle, they frequently use voice, a very natural means of communications. Unfortunately, when pilots communicate with the aircraft's onboard computers, they are often forced to wade through numerous levels of indenture to reach the appropriate command. However, new interface devices have lessened this problem. Touch sensitive overlays and voice controls are two means to achieve easy communications. Boeing's new 777 uses touch sensitive cursor control devices so that the Captain and First Officer can achieve easier interaction with the multiple AMLCD displays on the flight deck. As another example, some recent experiments with a connected word recognizer have shown that it is possible to use conversational commands and still achieve 99% recognition accuracy (Barry, Solz, Reising, and Williamson, 1994). The ultimate goal of the conversational commands is to emulate the interaction of the GR-1 and F-15E crews.

## 3.5 Present Summarized Information

Even though an efficient means of communication exists between the pilot and the EC, it does not mean that information always flows in a clear and concise way between them. In modern military aircraft

cockpits, pilots very often suffer from data overload, and with the inclusion of information from off-board sources, this overload problem may get worse. The designers of the EC can solve this problem by allowing the EC to present only summarized information. Icon based display formats, supplemented by text when necessary, are a very efficient way to achieve this goal. Steiner (1989) presented pilots with system status information in both text form and through icons. For very simple displays the pilots performed equally well with either type of display; however, as the displays became more complex, as measured in bits, the icon based format was clearly superior. The pilots comments also supported the fact that icon based formats were easier to interpret and gave better situational awareness.

## 4. CONCLUSION

In order to function effectively, the pilot and the EC must work together as a close-knit team. The ideal relationship between pilot and EC can be likened to that of the good manager and his staff. The pilot manager must be sufficiently aware of the work of the EC to be able to predict problems, but not so involved as to do its work. In other words, the pilot must be involved enough to be able to offer assistance when called upon, and yet must not micro-manage the EC and risk becoming overloaded and prevented from making strategic decisions. The good manager will know which staff members can be relied on to act without supervision, just as the pilot will form opinions of the strengths and weaknesses of the EC and decide when it can be given autonomy. As in the conventional management situation, the EC must maintain a knife-edge balance of providing sufficient data exchange without swamping the pilot system manager. By using the guidelines discussed in this paper, a good manager-staff relationship can be achieved within the H-EC Team, resulting in trust within the team. This internal trust will lead to an efficient and effective team. Others outside the team will see consistently correct performance, and will, in turn, gain trust in the team.

## 5. REFERENCES

Asimov, I. (1950) I, Robot. Fawcett Publications INC.

Barry ,T., Solz, T. , Reising, J. and Williamson, D. (1994) The use of word, phrase and intent accuracy as measures of connected speech recognition performance", To appear in Proceedings of the 38th Annual Meeting of the Human Factors Society, Santa Monica, CA: Human Factors Society.

Emerson, T.J. and Reising, J.M. (1990) The effect of an artificially intelligent computer on the cockpit design paradigm. In Proceedings of The Human - Electronic Crew: Is The Team Maturing? [Pp. 1 - 5] (Tech. Report WL-TR-92-3078). Wright-Patterson Air Force Base, OH: Cockpit Integration Directorate.

Haglebarger, D.W. and Thompson, R.A. (1983) Experiments in

teleterminal design. IEEE Spectrum, October, 40-44.

Howell, W.C. (1993) Engineering Psychology in a Changing World. In Porter & Rosenzweig (Eds.) Annual Review of Psychology (pp 231-263) Palo Alto, CA: Annual Reviews Inc.

Krobusek, R.D., Boys, R. M. And Palko, K.D. (1988) Levels of autonomy in a tactical electronic crewmwmber. In Proceedings of The Human - Electronic Crew: Can They Work Together? [Pp. 124 - 132] (Tech. Report WRDC-TR-89-7008). Wright-Patterson Air Force Base, OH: Cockpit Integration Directorate.

Muir, B.M. (1987) Trust between humans and machines, and the design of decision aids. International Journal of Man-Machine Studies, 27, 527-539.

Reising, J.M., Emerson, T,J. and Munns,(1993) R. C. Automation in Military Aircraft. In Proceedings of HCI International '93: 5th International Conference of Human-Computer Interaction, Orlando, Florida, August.

Steiner, B. and Camacho, M. (1989) Situation awareness: icons vs. alphanumerics. In Proceedings of the 33rd Annual Meeting of the Human Factors Society, Santa Monica, CA: Human Factors Society..

Turing, A. (1950) "Computing Machinery and Intelligence" in Feigenbaum, E. & Feldman, J. (Eds.) Computers and Thought, Malabar, Florida: Robert Krieger Publishing Company, 1981.

Wickens, C.D., (1992) Engineering Psychology and Human Performance, New York: HarperCollins Publishers Inc.

Wilson and Rutherford. (1989) Mental models: theory and application in human factors. Human Factors, 13, 617 - 634 .

# SESSION IV - CO-OPERATIVE INTERFACE DESIGN

## SYNOPSIS

In this section, six papers focus on methodologies and requirements for the design of HE-C interfaces. Paper 17 concerns an advisory system for making human-computer interfaces more reliable; Paper 18 describes an AI-based methodology and data base for task analysis and synthesis; Paper 19 concerns applications of a cognitive modelling design tool for human-system integration; Paper 20 reports an experimental study of automation invocation procedures; Paper 21 reports the evaluation of aircraft H-EC interface display formats; and Paper 22 considers the need for standards and engineering methods to promote trustworthy HE-C teamwork. *Paper 17* describes the development of the ARCHIE software kernal. The system monitors the operator and the environment, and provides advice about how the interface could be improved dynamically, in real-time. The aim is to help avoid and correct user errors, and to assist in management of operator workload. The system includes operator monitoring, optimisation of multi-modal inputs and controls, and plan recognition. Applications include interfaces for physically disabled people, as well as for ATC and cockpit environments. *Paper 18* describes a Human Factors Task Database, and considers the roles of functional programming and expert systems in human performance analyses. An example task history is described for an aircraft navigation scenario. The expert system aids task analysis by generating legal task combinations, by analysing potential workload problems, and by recommending design changes or task allocation options. *Paper 19* compares design concepts for integrating computers into cockpits and nuclear power plant control rooms. A NASA computer modelling tool (MIDAS) is described for computational modelling and prototyping of the cognitive implications of proposed interface designs, early enough for the latest ergonomics data to be useful. MIDAS outputs for a control room example are given. *Paper 20* reports an experimental study of the effects of different levels of pilot involvement in the procedure for automation invocation, i.e. the criteria for dynamic HE-C task allocation. Using the MINSTAR task battery, similar to MAT (see Paper 13), performance with four levels of automation was compared, ranging from complete pilot control to complete systems (i.e. EC) control. In the intermediate automation conditions, the shift from manual to automated status was cued by visual and aural warnings, or by changes in the tracking display presentation. The results are discussed in terms of the consistency of warning systems, fractionation of performance, and workload associated with invoked automation. Recommendations are made for adaptive allocation based upon a contextual model of the environmental constraints on human control strategies. *Paper 21* reports the design of "intuitive" HE-C display formats, and of associated controls, and their evaluation in man-in-the-loop flight simulations, with human emulation of an EC. The experimental cockpit simulated large area, flat panel displays, with a natural speech EC voice control system. Recommendations are made for head-up cues, voice control, EC prompts, simplified options presentation, and format removal. Consistent indications are given of preferred EC role allocations during high workload, with EC as systems and flight path manager, and the pilot as mission manager. *Paper 22* discusses the special nature of HE-C teamwork. By representing a knowledge system, rather than a program-level system, HE-C teamwork engenders specific requirements for design standards. From a review of the present state of HE-C technologies, and of associated standards, the author proposes that what is needed is a design process standard for HE-C systems operating at the knowledge level. The design and testing of HE-C teamwork is considered against threats to trustworthiness from competing perspectives. A framework is proposed for considering the level of maturity of HE-C teamwork. It is concluded that while at the current early stage of maturity, the way forward is a checklist approach to assessing teamwork, promoted under the umbrella of safety management.

# Improving the Reliability of Interactions in Human Computer Teamwork

Ian W Ricketts, Alistair Y Cairns & Alan F Newell
MicroCentre, Dept of Mathematics & Computer Science,
The University,
Dundee DD1 4HN
UK

Tel : +44 1382 23181   Fax: +44 1382 23435   Email: ricketts@uk.ac.dund.mcs

## Summary

With the increasing amount of computer power available in civilian flight decks it is becoming feasible to use some of this power for non-flight-critical systems. One area which could benefit greatly from some additional computer assistance is the pilot-machine interface. We describe the ARCHIE project. An investigation into making human-computer interfaces more robust and reliable. The initial target areas of this project are civilian glass cockpit flight decks and air traffic control stations.

## Introduction

"No matter how much has been done to design a workplace, its equipment and the worker's shift patterns to keep alertness at a maximum, there is no guarantee that people will not fall asleep on the job ... Imagine a future operator seated at the controls of a high-tech plant in the small hours of the night. She is well trained in sleep and alertness management, but the work is monotonous that night and her alertness starts to wane. Unobtrusive infrared eye-tracking technology or a light weight head-mounted monitor picks up the first signs. Without her noticing, the brightness of the illumination in her glare-free room is automatically increased, and the layer of air at head level is cooled. A stimulating aroma wafts into the room, while the computer flashes a low-priority but interesting task onto her screen. Instead of slipping into a zombie-like state, she is restored to full alertness, and is once again up to the task of coping with an emergency, should it arise.
...When such technology becomes a standard feature of 24-hour operations, our society will finally have reversed the tendency that leads to so many industrial breakdowns and disasters"[1].

Our aim is to develop a transparent software kernel which runs in parallel with the flight software, monitoring the operator and the cockpit environment, and providing pertinent timely advice. There have been a few attempts at this sort of assistance on the flight deck before, but the ARCHIE approach is unique in that it is a generic system, not tied specifically to any one application (such as flight management software, air traffic control software, etc.).

ARCHIE is a collaborative project with kernel development work being shared between Dundee and Computer Resources International (Denmark). The other partners are GEC-Marconi Avionics (UK), Bertin & CIE (France), Sofreavia (France) and the UK Civil Aviation Authority.

The goal of the ARCHIE project is to create a software kernel which, when run in parallel with an existing application, will increase the reliability of the human computer interaction between the application and it user. The kernel monitors the operator, the environment and the current state of the application and, on the basis of this information offers advice about how the interface could be improved dynamically in real-time. It also offers the ability to use alternative input modalities for non-flight-critical functions (such as radio tuning) and attempts to predict (and thus help avoid) user errors. The system is designed to run on a heterogeneous distributed network of computers, with the application software running under the UNIX operating system in an X windows environment. Typically, the bulk of the kernel runs on UNIX workstations, with the additional input device processing being done on separate, smaller machines, such as IBM personal computers, all communicating over an Ethernet network. Although the environment described was used for the development work, the ARCHIE system is not tied to any specific operating system, architecture or windowing system.

The kernel is purely an advisory system, offering advice to the application which is free to act on it or ignore it as it sees fit. Although little or no modification of the application is needed to use the ARCHIE kernel in its most basic way, the

This statement means that every ten seconds, the code in the body of the statement will be executed, based on the statistics for the previous five seconds. This code will produce a report (which is logged to a file) consisting of the phrase "Head Data Report" and the mean value of the readings from the device OpDist (for the last five seconds). Then, if the standard deviation of the readings from OpDist, again over the last five seconds, is sufficiently low (indicating a lack of head movement and perhaps fatigue on the part of the operator), a recommendation to make all output come from a speech synthesiser, in an attempt to ensure that the (possibly inattentive) operator still receives information from the kernel. As always, it is up to the application to act on this advice or to ignore it.

**Multimodality**

The second additional feature that an ARCHIE system provides is the ability to add alternative modes of input and control for a computer system. It is generally accepted that the inclusion of such additional modes of input increases the efficiency, and hence the reliability, of a human computer interface [2]. The ARCHIE multimodal system takes input from a variety of different devices (e.g. gesture, speech, eye gaze) and combines and modifies them, producing output which is sent to the application. This output is made to look as if it had come from the existing input devices, such as mouse and keyboard. The main advantage of this device independent approach is that the application does not have to be changed (other then the addition of a few lines of code) before it can be used with the new devices, since it still believes that it is receiving input only from the mouse and keyboard.

The system has two parts, one dealing with the translation of input into internal symbols and the other with the generation of output from these symbols. This split allows the system to be easily ported to other windowing environments by simply modifying the second part, which is dependent on the peculiarities of the current environment. It also allows, by providing an intermediate stage, input data to be fused so that several inputs generate one internal symbol (such as waving and saying "Hello").

As with the monitoring system described above, the multimodal system is programmed by the application developer using a supplied language, the MIDL (Multimodal Interaction Definition Language). Programs written in this language have a very simple format, allowing the multimodal interaction to be specified quickly and intuitively by a non-expert. We see this as an important feature, allowing developers to "try it and see".

The language essentially has three parts: the definition of input data, the specification of how inputs are mapped to internal symbols and how these symbols are used to generate output. A simple example of this is given below:

```
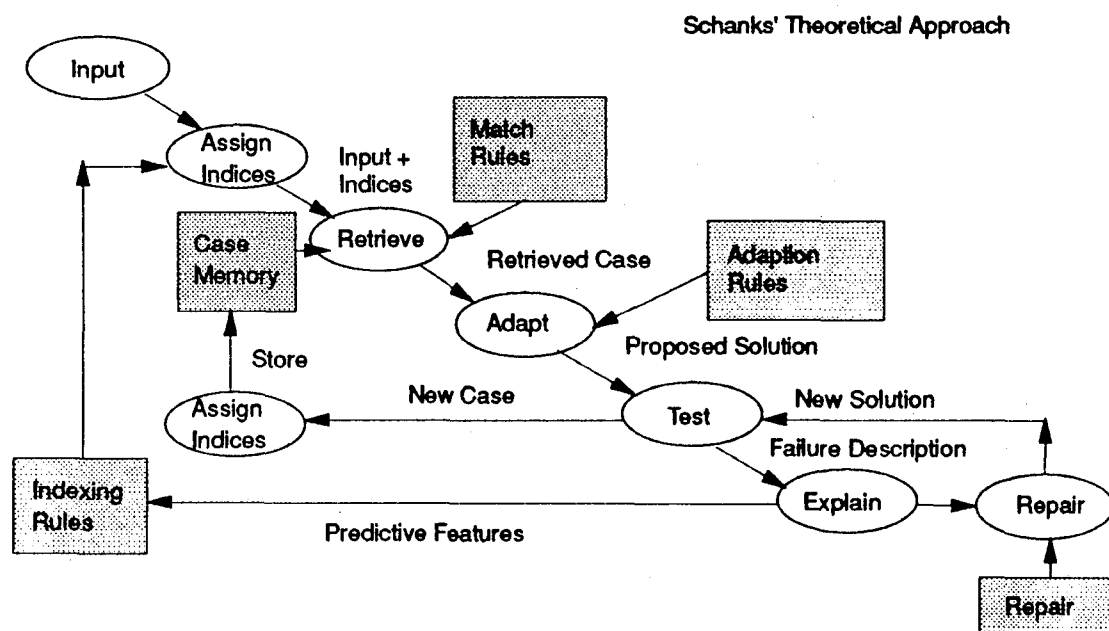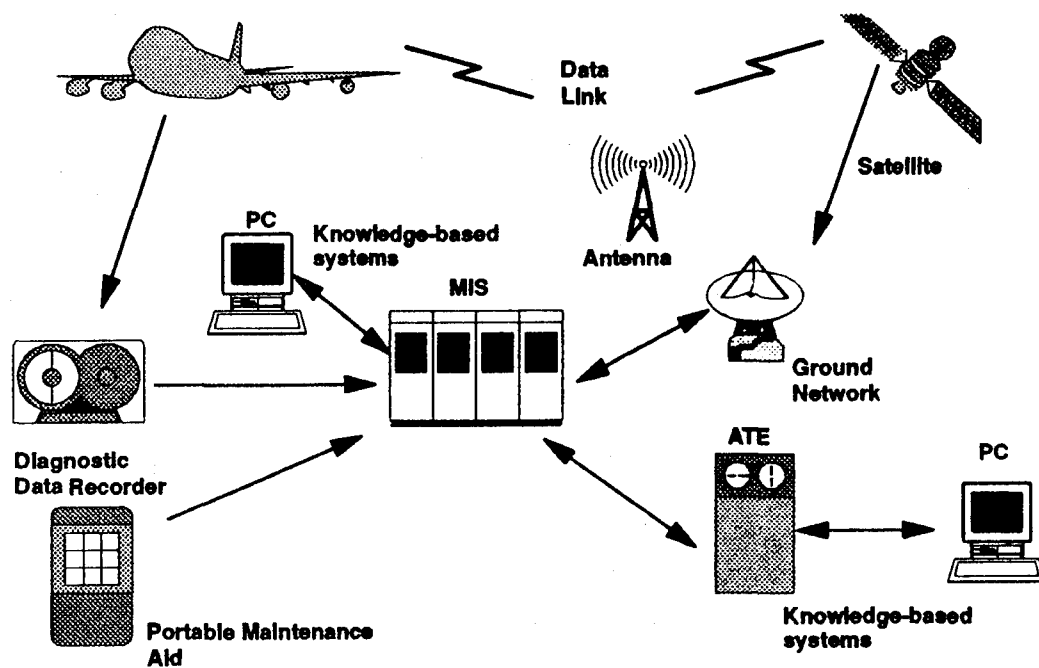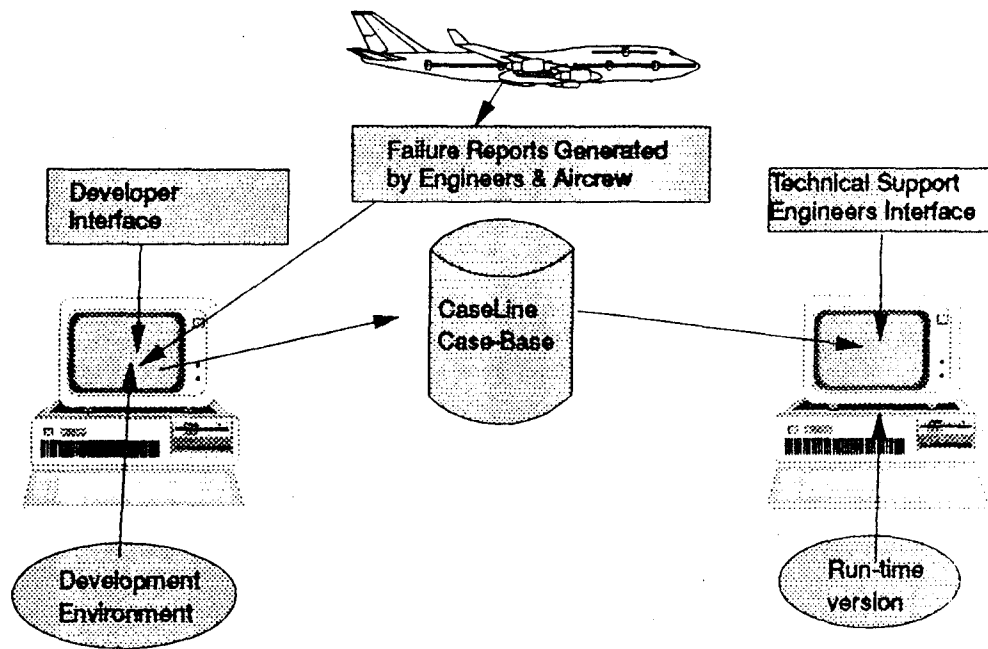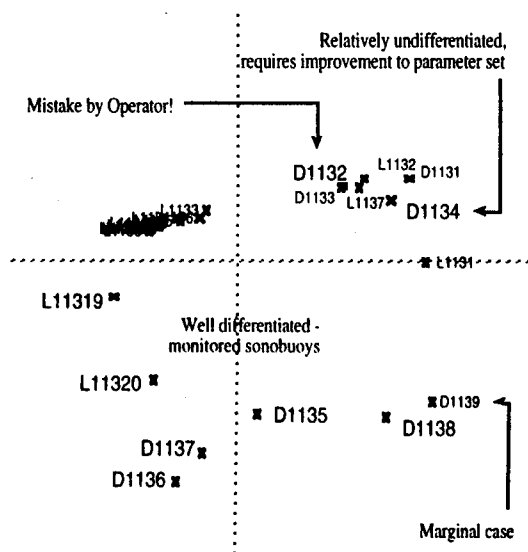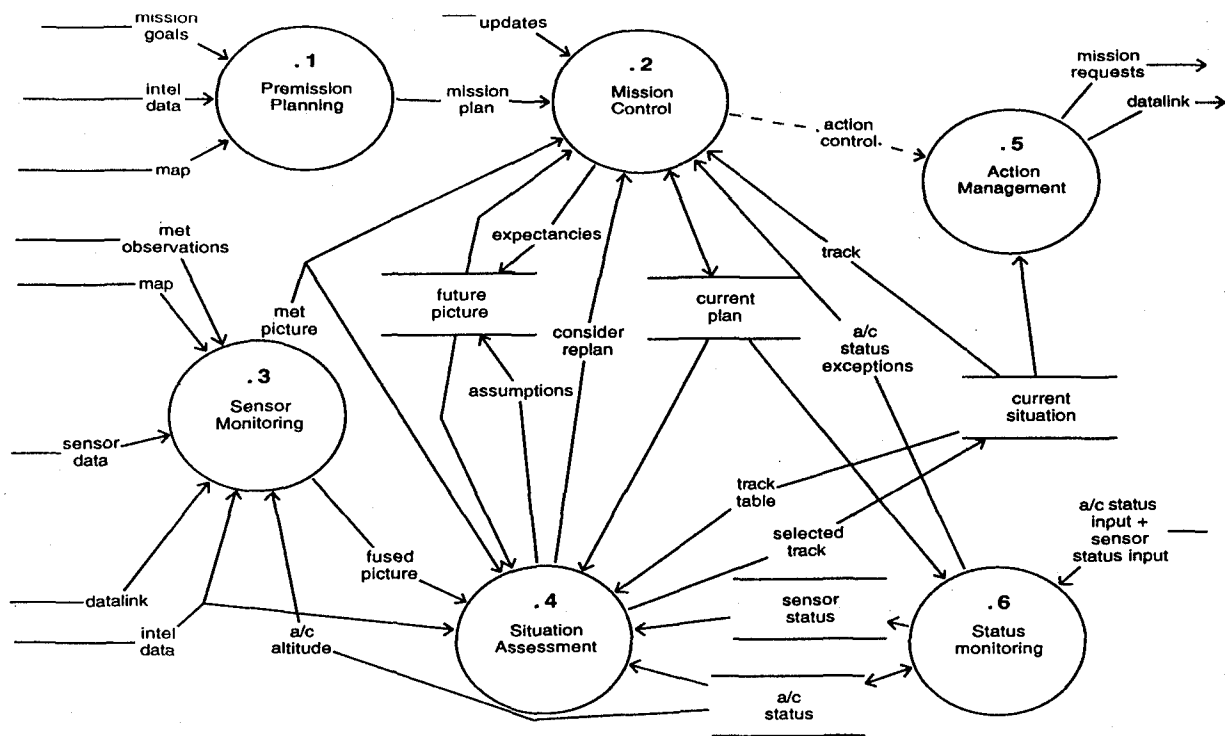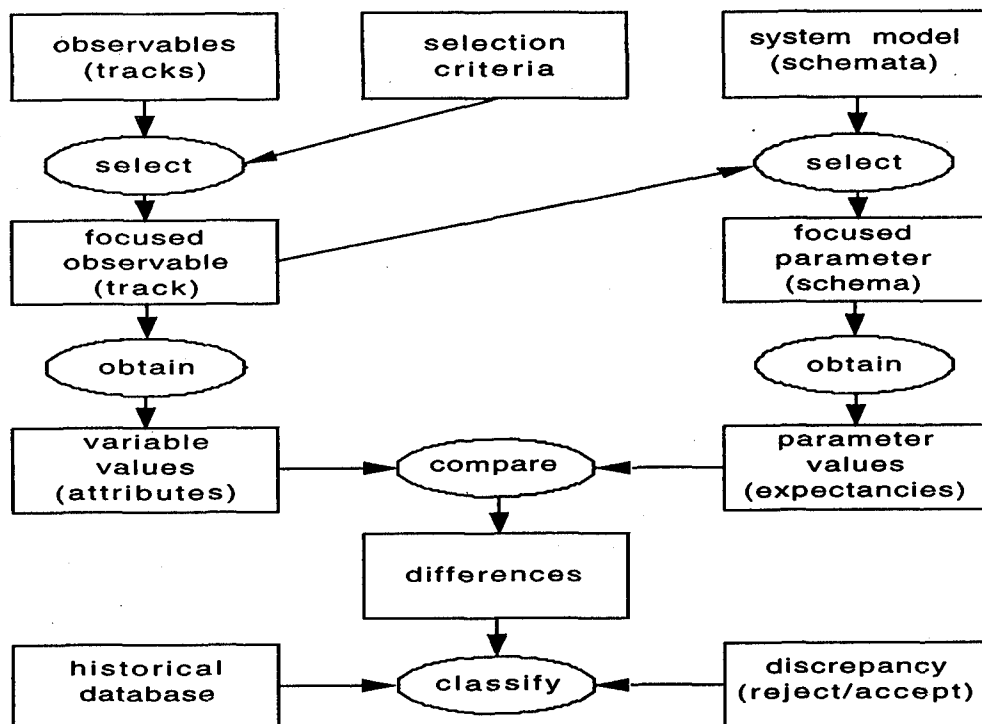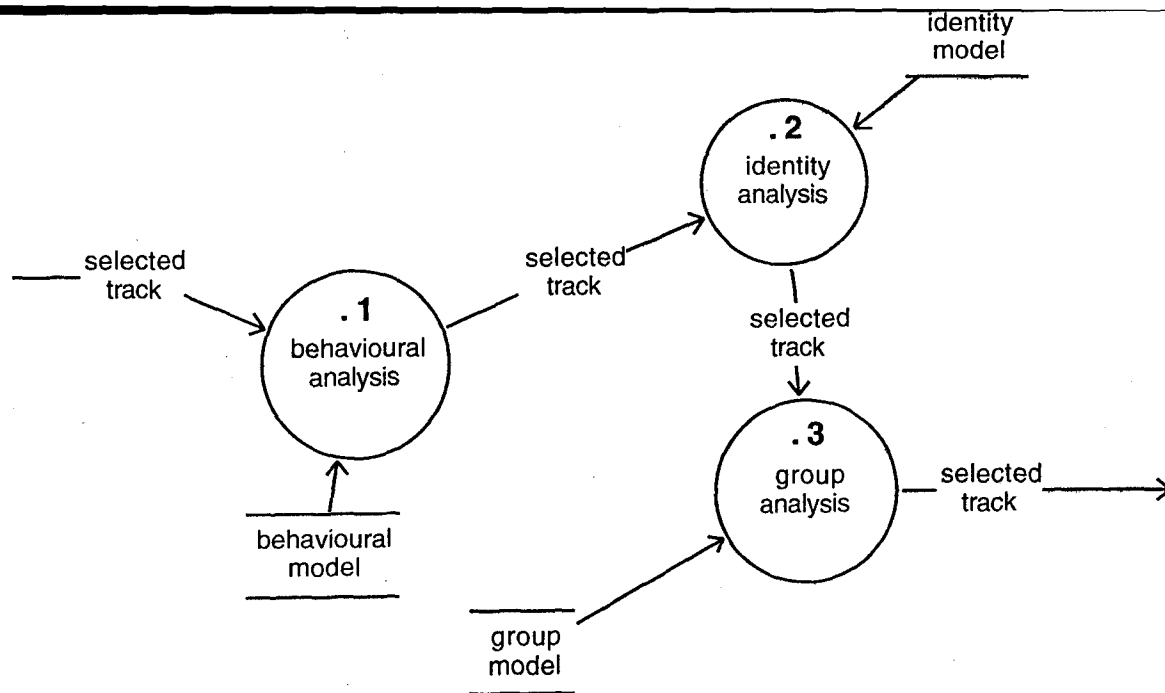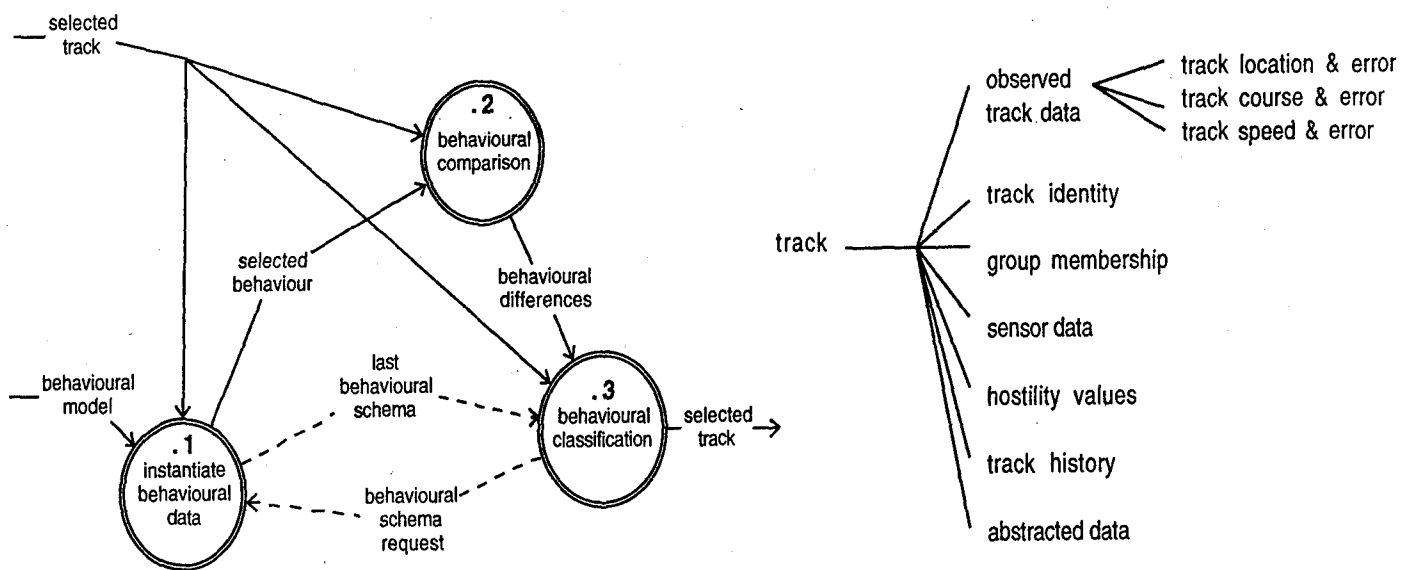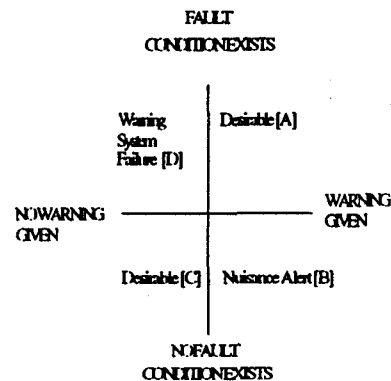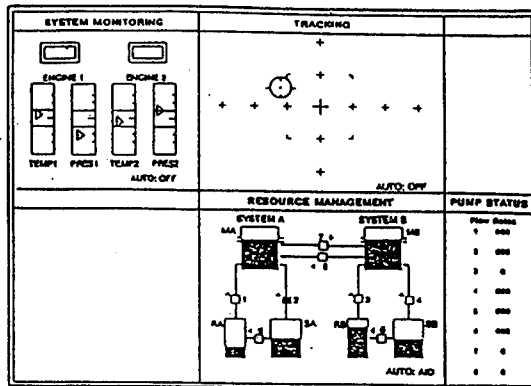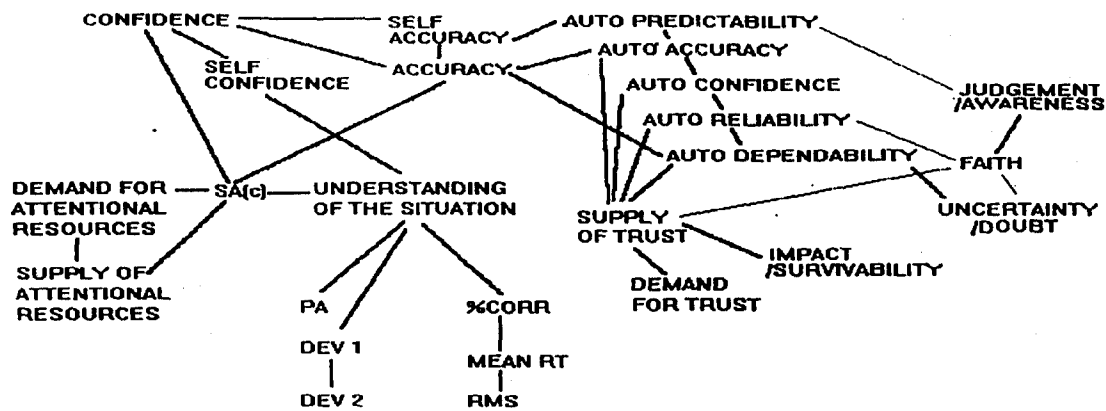define input Wave 75;
define input TalkHello 45;

define symbol WaveTalkSymbol
{
device glove Wave;
device speech TalkHello;
check device SwitchOne 1;
within 1500;
}
```

This defines two inputs, *Wave* and *TalkHello*, specifying the values (sent by the additional devices) which represent them. The *define symbol* statement defines the conditions for the generation of the internal symbol *WaveTalkSymbol*. These conditions are that it receives the input *Wave* from the device called *glove* and the input *TalkHello* from the device called *speech* within a time interval of no more than 1500 milliseconds between the first and last inputs. When these are received (in any order), it checks to ensure that the last value supplied by the device called *SwitchOne* was one. If this is so, the symbol is generated. Typically, several of these types of statements are present in a MIDL program.

The last set of statements defines how the internal symbols map to the generated output. For example:

```
define effect
{
trigger SymbolOne then SymbolTwo
        then SymbolThree time 500;
output click 1 to DisplayWidget;

if (SymbolTwo[1] = = empty)
        output keysequence "hello" to
        (eye.x, eye.y);
else
        output click 2 to DisplayWidget;
}
```

This defines the top-level plan, called **end**, to consist of either the plan **hunt, robbery** or **cash_cheque**. In turn, it defines each of these plans to consist of their own individual sub-plans, until everything is reduced to atomic actions. For example, **hunt** has a precondition (which must be satisfied before the plan can be considered as a potential one the agent is following) that the variable **hunting_season** is set to true, i.e. you cannot hunt out of season. Assuming that this condition holds, the plan consists of the sub-plan **prepare-hunt** which is defined further on in the program. Conditional statements are allowed in plans, as can be seen from the **robbery** plan. If the variable **agent** associated with the plan is equal to either "Joe" or "Jack" and **has_a_criminal_record** is set to true, then the plan has a sub-plan called **advanced** otherwise it has a sub-plan **beginner**. The reduction from the top-level plan **end** down to the individual actions which are used in the recognition process must be completely detailed in the PDL program.

**In the Cockpit**
Although the ARCHIE kernel is designed to be generic and not tied to any one application, it must be used with an application in order to demonstrate its capabilities. We have chosen two application areas in which to test the kernel, the next generation of air traffic control station and the civilian glass cockpit flight deck. These were chosen for a number of reasons. They are both areas where humans interact with computers, they both involve significant workload and stress for the operators and the consequences of an operator error can potentially be very serious. In the remainder of this section we will discuss the potential benefits of ARCHIE in a future flight deck.

All flight deck operations, particularly in the civil field, follow a rigidly defined set of procedures, specifying the actions to be carried out, the order of execution and, where applicable, any required pre-conditions. These instructions may be used directly to instantiate the plan recognition part of the ARCHIE kernel, which can then be used to monitor the pilot's action sequence, in conjunction with aircraft state parameters, to detect deviations from the permitted procedures.

The ADL capabilities of the ARCHIE kernel would be of use in two respects: monitoring of the operator's workload/physiological status, and monitoring of the overall aircraft state. In the former case ARCHIE could detect, for example, a low level of pilot alertness and tailor the environment and/or information displays to counteract this. Options may be to reduce the cockpit temperature slightly or to invoke a visual or audible attention-grabber. By monitoring the aircraft state, the ARCHIE system will be capable of detecting (and providing timely warnings about) a range of subtle trends, not detectable by current central warning systems, which could lead to a dangerous situation.

By monitoring the status of the cockpit environment, the ARCHIE kernel would be able to deduce the optimum means of imparting information to the crew - for example by adopting visual rather than audible warnings if the ambient noise level exceeds a pre-defined threshold. The MIDL capability of the ARCHIE kernel provides a means of integrating alternative command-input devices into existing cockpit systems. While the multimodal input capabilities offered by the kernel are not currently envisaged as being particularly useful in the civilian cockpit application, they may be of great use in a high-workload military cockpit environment to provide a moment-to-moment optimisation of aircrew command-input mechanisms. For example, in situations where manual activation of a switch is not possible (due to workload or g-loading) the ARCHIE system could provide the capability or activating the switch based upon the input from an eye-gaze sensor.

One of the ARCHIE consortium partners (GEC-Marconi Avionics Ltd) is currently developing a demonstrator system to illustrate some of the concepts of ARCHIE in the civil flight deck. The demonstrator has been developed on commercially available computing equipment and provides an instrument display based upon that of the Boeing 757 airliner. It incorporates an ARCHIE kernel which is instantiated to monitor the actions of the operator against the known sequence of actions required to complete the various flight procedures. The plan library developed for this purpose was derived directly from the procedures specified in the official aircraft operations manual. The demonstrator also makes use of the ADL facilities of the kernel to monitor the aircraft state for detection of dangerous trends and uses both audible and visual means of providing advice/warnings to the operator.

The main scenario selected to demonstrate ARCHIE in the civil cockpit is the approach-to-land procedure (including a potential go-around).

# Towards an Expert System for the Analysis of Computer Aided Human Performance

A. N. Clark, G. L. Greatorex & A. G. Hill
Computer Systems Division, GEC-Marconi Research Ltd.,
Great Baddow, Chelmsford, Essex, CM2 8HN, Great Britain
e-mail: anc@gec-mrc.co.uk, tel.: (0245) 473331

## ABSTRACT

This paper summarises work which has been conducted at the GEC-Marconi Research Centre as part of a collaboration between the Human Factors and Knowledge Based Systems Groups. The work is aimed at increasing the cost effectiveness of human factors engineering analyses by introducing computer and expert systems. This work is particularly relevant for aircraft developments where human performance can have a significant impact on mission success. Specifically, a Human Factors Task Database to facilitate storage, retrieval, and reuse of task analysis data is described. Furthermore, the representation of task and scenario attributes in a computer readable form is discussed. This would facilitate the use of expert systems to analyze design and task allocation options in order to propose solutions which promote effective human performance.

## INTRODUCTION

Mechanisms for the cost effective assessment of aircrew performance (including decision making) in complex aircraft systems can greatly enhance design for operational effectiveness. A key component in any such assessment is task analysis, because human performance is highly dependent on the context in which tasks are performed. Unfortunately, task analysis is a time consuming and resource demanding task. In addition, subsequent human performance evaluation is often error prone because of the large number of variables which can affect human performance. This area is ripe for computer based support in the form of databases to facilitate storage, retrieval and reuse of task information and also expert systems for analyzing these task descriptions.

This paper describes ongoing research at the GEC-Marconi Research Centre, which is a result of a collaboration between the Knowledge Based Systems and Human Factors Groups within the Computer Systems Division. It addresses the fundamental issues which are involved in precisely describing a *task* with a view to providing a basis for computer aided support of the task and human performance analyses which are performed by a Human Factors expert.

The paper is presented in five sections including a description of the Human Factors Task Database (HFTD), the envisioned role of expert systems in human performance analyses, an explanation of functional programming including the task elements which must be defined to allow the use of expert systems, an example, and a conclusion. We intend that the basic description of task elements and human performance analysis which is given, convinces the reader that this technology provides the basis for a broad range of computer aided analyses.

## HUMAN FACTORS TASK DATABASE

The HFTD has been developed to increase the cost effectiveness of task analysis, to facilitate the reuse of task information, and to provide a structured framework for data collection and analysis. Central to the approach are the principles of hierarchical task analysis, where tasks are decomposed into their component subtasks to facilitate analysis. [1] A wide variety of data can be entered into the database for each task. For example, descriptive information such as the purpose, action, potential error modes, information requirements, skill requirements, and decisions associated with a task can be entered. In addition, more quantitative information including workload, time to perform, error rates, level of difficulty or importance can be stored. Furthermore, data about the conditions under which quantitative data were assigned, or the reliability of data can be included.

The HFTD also allows for the construction of scenarios from a combination of tasks or task hierarchies. This allows for a time sequence of tasks to be built up, and also for a plan under which tasks are performed to be specified. For example, whether two tasks will be performed in sequence or in parallel, or the conditions which must be met prior

to a task starting. At the scenario level, conditions specific to the scenario which affect task performance, such as fatigue, time of day, or other tasks which must be performed in parallel must be considered. The HFTD allows for a library of these factors, and their impact on task performance to be collected and utilised.

The primary benefit of this database arrangement is that for a given scenario or task hierarchy, any combination of task data (*e.g.* time to perform, workload, information requirements) can be retrieved for further analysis or for a report.

## ROLE OF EXPERT SYSTEMS IN HUMAN PERFORMANCE ANALYSIS

A human factors expert who is faced with trying to predict human performance in a complex system like an aircraft cockpit has a difficult challenge indeed. Expert systems can greatly facilitate this process. [2] At the most basic level, they can take data from the HFTD in terms of the time, workload, human error, or human resources associated with a scenario, and compare them with operational requirements. This assessment can be taken a step farther and "what if" questions can be asked about function allocation between the aircrew and the onboard computer systems, task allocation between crewmembers, and task scheduling options.

More complex human reliability analyses including decision making performance could also be conducted. Expert systems could sort through large tasks and scenarios to ensure the correct information is available for each task, or checks could be done to determine if information is presented consistently for all common tasks. The number of steps required to access information could be reviewed and assessed against time to perform task or resources used. At a more complex level, the quality of information can be manipulated to assess changes in performance, or the options for sharing of information between crewmembers can be evaluated. Furthermore, decision making performance could be adjusted by performance shaping factors (*e.g.* fatigue, motivation, time of day) or underlying workload. [3]

## FUNCTIONAL PROGRAMMING

In order to use expert systems during task and human performance analyses, relevant task attributes must be represented in a computer analyzable form. Of course, the specific attributes which are modelled

will be highly dependent on the type of analysis being conducted. Therefore, it is important that the computer representation of the task is flexible enough to lend itself to "what if" type questions and to the wide variety of analyses (*e.g.* timeline, workload, human reliability, decision making, *etc.*) which may be performed. For this reason, we have chosen functional programming because it's notations provide a particularly simple and concise representation which is amenable to varied mathematical analyses in addition to execution as a conventional programming language. [4] The following sections will provide some examples of the types of task attributes which need to be represented in a mathematical and/or logical notation to enable the use of functional programming in the analysis of human tasks. Generally, these attributes fall into three basic categories: attributes of the task, attributes of the agents (whether human or machine) that perform the tasks, and attributes of the scenario or the way in which tasks are executed. This information will be used by the expert system in two primary ways. First, in order to give meaning to the task execution (*i.e.* when constructing a computer analyzable task scenario), and in conducting one of a number of task or human performance analyses.

### Task Attributes

The minimum information which must be provided to the expert system about a task is as follows:

a) Inputs (displayed information, system state, *etc.*)
b) Activities performed during the task (*e.g.* controls activated)
c) Outputs (change in system state, *etc.*).

In addition, depending on the type and sophistication of the intended analysis, more detail can be added. For example:

a) Error rates (including decision making)
b) Time to perform
c) Information requirements (availability, quality, and priority of data could be quantified)
d) Decisions required during the task. This could be expanded to describe the importance of the decisions, the information required for each decision, the quality of information provided, the number of steps required to access the information, the

probability that the decision will be made with the help of other agents (other crewmembers or the computer) competence required to make the decision, *etc*.

e) Workload
f) Skill requirements (trained performance standards could be specified)
g) Equipment requirements
h) Task Priority

## Agents

Each task is performed by a collection of *agents*, in a flying example, the task is performed by a pilot, possibly in conjunction with other agents such as a co-pilot, a navigator, or an onboard computer system. Agents work concurrently to achieve a collection of tasks, interacting where necessary to achieve an overall goal. Each agent proceeds by observing and acting upon attributes of the task, for example a pilot flies the plane by moving switches, pushing levers, requesting information, reading head up displays, *etc*. Agents interact by sending each other messages, for example the pilot may request the next waypoint from the navigator by sending a "*next waypoint*" message which will produce the reply message "*next waypoint at X*".

Agents also have attributes associated with them. In many cases, these are thresholds which define their performance capabilities. For example, skills they have acquired, their ability to perform specific tasks within a given time or accuracy, working memory capacity, information processing limits, or workload thresholds. In the case of workload, the expert system can assume that above this threshold there is no guarantee that the agent will successfully perform a requested task. At any time, an agent is performing 0 or more tasks. Each atomic task is associated with an effort value which describes how the task affects the agent in terms of workload. If any step performed by an agent would cause the aggregate of the effort values to exceed the workload limit for that agent then that step can fail. Failure of a task step can cause an attribute not to be updated, failure to request another agent to perform a task, failure to make a proper choice between two subtasks or failure to produce a requested outcome.

Initially, at the start of scenario task execution, all agents are inactive. Then, as the overall task progresses, agents are requested (by the expert system) to perform individual tasks, each of which may initiate sub-tasks. Agents may interact and send each other messages in order to co-operate

when achieving tasks. When an agent receives a message it will attempt to initiate the request immediately. Any task performed by an agent may be delayed, or may fail in whole or in part because of workload restrictions.

A verbal specification of an agent is provided below. This could be expressed in a mathematical form for use by a functional program.

An *agent* has
    a collection of attributes; and
    a workload limit; and
    a collection of named tasks to perform
        where a *task* is executed as either
        a choice between two subtasks; or
        a sequenced pair of subtasks; or
        a parallel pair of subtasks; or
        an attribute update; or
        a request for another agent to perform a
        named task

## Task Scenarios/Execution

When we think of a *scenario* it is usual to envisage a series of actions which are performed, using a collection of tools and resources, in order to achieve a particular goal. For example, flying an aircraft along a preset course from waypoint to waypoint to reach a final destination. Along the way there may be tasks to be performed before the overall goal is achieved and it may not always be clear in which order the tasks have to be performed. For example, whilst flying the aircraft, the pilot is responsible for broadcasting the aircraft position at preset times and maintaining the correct altitude along with other tasks. Often it will not matter which order these tasks are performed so long as the aircraft gets to the intended destination at the required time.

An expert system written in functional programming language can be set up to execute the tasks in a scenario either in a prescribed sequence, randomly, or according to rules such as the probability of performing various tasks. In fact, this execution can be varied to produce different properties of the task, or the manner in which it is accomplished, in order to observe the effect which these modifications have on performance against scenario objectives.

Each execution of a scenario produces a *task history* which is the sequence of steps from the start of the scenario to the end (which may or may not achieve the desired goal). The task history contains all of the information which is used at every stage of the

scenario, showing, for example, messages which are sent between agents, the calculations which lead to decisions being made and actions which affect the task attributes. The task history for the aircraft example is like a perfect flight recorder which enables all of the steps of the flight to be played back in minute detail.

*Verbal Specification of a Task History*

A *task history* is a sequence of *steps* producing an *outcome*
    **where** examples of *steps* include
        initiation of one or more tasks; or
        performing a decision; or
        producing an outcome; or
        updating an attribute; or
        requesting another agent to perform a task; or
        failing to perform a task

## Task and Human Performance Analysis

Given a task history, a Human Factors expert will be able to analyze the steps which were performed in order to identify evidence of bad task or system design. Specifically, this will relate to areas where human performance targets are not achieved. For example, certain targets may not be achieved because the pilot tries to perform too many activities at once. In this sense, the Human Factors expert is debugging the trace of a buggy task execution – the analysis may be performed because the overall goal is not achieved and the expert is trying to trace back to the place in the history at which the task starts to go wrong; alternatively, the expert may be unhappy about the overall performance of some of the agents in the scenario and wishes to analyze the task history in order to identify areas where the agents can be made to perform better.

Furthermore, it is possible to identify task histories which lead to performance problems, for example, due to invalid information being available because one agent fails (because of overwork) to respond with up-to-date information when requested by another agent. Such a situation would occur when at least one of the task histories associated with a scenario can be shown to fail due to an agent being overworked and thereby preventing a second agent from receiving up-to-date information which it requires in order to make a decision.

If each agent interacts with the task attributes by making decisions based on its understanding of the situation at the time. The validity of the agent's

decisions depends (among other things) on whether or not the information which it uses to make the decision is up-to-date or not. Many of the task attributes will change frequently and it is important that all agents which make decisions based upon these attributes keep their knowledge of the values current by monitoring the attributes at frequent intervals. For example, if the pilot aims to fly at a constant altitude then it is important that the altimeter is monitored frequently otherwise the decisions which are made involving increasing or decreasing the altitude are likely to be invalid.

In addition to detecting potential problems with task executions, expert systems allow observation of what happens when attributes are modified. Going a step further, expert systems technology, in the form of heuristic rules, can be used to guide a program through a collection of task modifiers in order to identify and remove the source of a problem.

## EXAMPLE SCENARIO

We believe that the proposed task, agent, and scenario descriptions will allow a sophisticated degree of task and human performance analysis. To illustrate its use, we present a simple example in this paper based on agent workload. Given the task execution mechanism described above, we will show how this analysis can be captured by constructing a computer readable scenario definition which is inspected for optimum agent workload levels.



**Figure 1.** Example Scenario

Figure 1 portrays an example scenario of an aircraft flying between waypoints. There are three agent nodes labelled *Pilot*, *Navigator* and *Cockpit*, each of which is connected to task nodes for each of the

named tasks which the corresponding agent can perform. Each task node is connected to a collection of nodes which in turn are connected to task nodes which can be invoked in series or parallel.

The scenario is composed of three agents, the cockpit, the pilot and the navigator. The cockpit agent is responsible for maintaining the information relating to the speed, altitude, direction and joystick position. The cockpit provides instrument readings for the speed and altitude which are available to the pilot and navigator. The cockpit also provides a task which is used to update the joystick position. The main task of the cockpit is to regularly update the speed, altitude and direction values using the position of the joystick. The workload of the cockpit is unlimited.

The pilot agent is responsible for keeping the aircraft on course by monitoring the instrument dials and updating the position of the joystick through the cockpit agent. The pilot is also responsible for requesting the next waypoint from the navigator when the current waypoint is being approached. At predetermined waypoints, the pilot must broadcast the aircraft position. The pilot's workload limit and the effort for each task, are predetermined values.

The navigator agent is responsible for maintaining an up to date position of the aircraft and for responding to the pilot's request for the next waypoint.

A verbal specification of the cockpit agent which could be expressed in a mathematical form for an expert system is provided below. Since it is not a human resource it has unlimited workload capacity and maintains the attributes which describe the speed *etc.* of the aircraft. We assume that there is some notion of global time whereby any changes in the controls are implemented causing the speed *etc.* of the aircraft to be updated. The cockpit offers a task to update the position of the joystick, which is how the pilot flies the aircraft.

A *cockpit* is an agent with unlimited workload
   with attributes for altitude, speed, direction and
   joystick position; and
   with tasks to perform the following
      get the current altitude; and
      get the current speed; and
      get the current direction; and
      update the position of the joystick; and
      fly the aircraft which involves a condition
         if it is time to update the displays

         then perform the following tasks in
         sequence
            update the altitude, speed and direction
            from the joystick position; then request
            the cockpit to continue to fly the aircraft
         else request the cockpit to continue to fly
         the aircraft

Similarly, a specification of the pilot agent is presented below. During the flight the pilot will monitor the instrument displays in order to check whether the joystick needs moving, whether a new waypoint is required and whether the aircraft position needs to be broadcast. The pilot is a client to both the cockpit and the navigator. Each of the tasks performed by the pilot has a predetermined effort.

A *pilot* is an agent with a maximum workload of ωp
   with attributes for altitude, speed, direction and the
   current waypoint
   with tasks to perform the following
      read the cockpit displays for altitude, speed and
      direction; and
      check the current position which involves the
      following sequence of tasks
         concurrently, check all instrument dials; then
         depending on whether the aircraft is off
         course
            move the joystick accordingly; then
            continue to monitor the position;
            otherwise
         just continue to monitor the position
      depending upon whether the aircraft is
      approaching the waypoint
         request the next waypoint from the navigator;
         then
         update the joystick accordingly; then
         continue to monitor the waypoint; otherwise
      just continue to monitor the waypoint
      depending on whether the current waypoint is
      prearranged
         request the current position from the
         navigator; then
         broadcast aircraft position information; then
         continue to test whether the position should
         be broadcast; otherwise
      just continue to test whether the position
      should be broadcast
      start the flight by
         initializing the attributes; then
         throughout the flight check position, the
         waypoint and whether to broadcast position

The specification of the navigator agent (not shown) follows the same lines as the pilot and cockpit and provides tasks to get the current position and the next waypoint. The navigator is a client of the cockpit.

In this case, the scenario definition is a collection of *agent nodes* each of which contains an agent and has an edge leading to a *task node* for each of the named tasks which the agent can perform. A task node contains the name of the task and the maximum effort (in terms of workload) which the task is likely to take. Each task node represents the possible executions of a task, since we are interested in workload, a task node is also associated with all possible combinations of server tasks which it can request in parallel. Each task node has a collection of edges, each of which leads to a set of task nodes.

A particular scenario can be *overloaded* if, given an initial collection of tasks, it is possible to exceed the workload limit of an agent node by constructing a collection of paths (task executions) as follows: for each current task node, select a set of server task nodes which it activates in parallel, add these nodes to the current collection and repeat. The workload of an agent is exceeded when there is a path which activates enough of its tasks to cause it to potentially fail. If there is a path which loops through an agent's task, then without further information to the contrary this will cause the agent's workload limit to be automatically exceeded.

As shown above, there is a possibility of overload because each of the tasks *Next Waypoint*, *Position* and *Check Position* can be performed by the pilot in parallel through the *Start Flight* task. So when $\omega p < \omega 9$ where $\omega 9 = \omega 6 + \omega 7 + \omega 8$ there is a possibility that the pilot will be overloaded.

In order to prevent overloading there must be some mechanism for preventing certain tasks from being performed in parallel. An expert system can be used to investigate simple extensions to scenario descriptions which specify the legal combinations of tasks. Techniques such as predicate path expressions can help determine which execution (task history) or set of executions (set of task histories) result in the best performance (in this case, acceptable workload levels). [5] A basic example of a program for this process is provided below.

For a system of agents A, with attribute descriptions V, performing a task T:

If T is a decision D leading to one of tasks $T_1$ or $T_2$,
    then the best execution (set of task histories) depends upon the result of decision D with respect to attributes V.
        If the decision is true then the best execution is defined by A, with attributes V, performing task $T_1$;
    else it is A with attributes V performing task $T_2$; else
If T is a sequence of tasks $T_1$ then $T_2$, then the result is the definition of A with V performing $T_1$ producing new attributes V' followed by the definition of A with V' performing $T_2$; else
If T is a parallel combination of tasks $T_1$ and $T_2$ then the result is the set of definitions constructed by permuting all possible combinations of performing each atomic step of $T_1$ and $T_2$; else
If T causes the agent to overload then the result is failure of T; else
If T is an attribute update then the result is V' where V' is the same as V except that the particular attribute has been updated
If T is a request for agent A to perform a task named n, then the definition is the definition of A with attributes V performing task T'
        **where T' is the task named n of agent A.**

This basic example could be expanded to encompass more complex analyses and more sophisticated scenario descriptions. Furthermore, this type of expert system, which helps predict human performance, can be tested against real world applications and modified/improved to provide increasingly accurate representations of areas within system design and task allocation which negatively impact human performance.

## CONCLUSION

This paper has addressed computer-based tools which can aid the analysis of human performance in complex systems. In the first instance, a Human Factors Task Database can greatly facilitate the storage, retrieval, and reuse of task analysis data. In addition, recording this data in an electronic medium allows for its further manipulation into a computer readable form which can be analysed by an expert system. The basis of the expert analysis is a precise specification of task and scenario attributes using a simple functional notation which has been extended with agent expressions. The resulting task notation has been given a semantics, defining precisely how tasks are performed. This

forms the basis for a Human Factors tool which can be used to prototype task specifications where the tasks are performed by interacting agents. Using knowledge from a Human Factors expert, the tool can analyze the task specifications in order to identify potential problems, or to recommend system design or task allocation options which enhance human performance. Functional notation provides a useful mechanism for representing tasks and scenarios. Since it is easy to manipulate and expand, it has real potential for and improvement based on actual performance data. It also lends itself to adaptation for many human factors analyses such as workload analysis, timeline analysis, assessments of human decision making, and human reliability analysis.

## REFERENCES

1. Kirwan B., and Ainsworth L.K., *A Guide to Task Analysis,* London, Taylor and Francis, 1992. (ISBN 0 7484 0058 3).

2. Benyon-Davies, P., *Knowledge Engineering for Information Systems*. McGraw-Hill International Series in Software Engineering, 1993.

3. Chainaik, E., & McDermott, D., *Introduction to Artificial Intelligence*. Addison-Wesley, 1985.

4. Field, A.J., and Harrison, P.G., *Functional Programming,* Wokingham, Addison-Wesley, 1988. (ISBN 0 201 19249 7).

5. Andler S., *Predicate path expressions,* In Principles of Programming Languages, 1979.

# MIDAS IN THE CONTROL ROOM:
## APPLYING A FLIGHT DECK COGNITIVE MODELING DESIGN TOOL TO ANOTHER DOMAIN

Douglas G. Hoecker & Emilie M. Roth
Westinghouse Science & Technology Center
Pittsburgh PA 15235
hoecker@cognac.pgh.wec.com

## SUMMARY

This paper briefly outlines the current state of design concepts for large-scale introduction of computers to advanced control rooms for nuclear power plants, and compares this with the similar technological shift that can be observed in the cockpit. Although the two domains have some distinct differences, the similarities in terms of implications for advanced human-computer interaction (HCI) are perhaps more striking. With the similarities in mind, we have begun to adapt a computer modeling tool, the Man-machine Integrated Design and Analysis System (MIDAS) that was originally developed to assist in the design of advanced cockpits. The adaptation effort is resulting in new functionalities to the MIDAS system that will generically improve its capacity to support the design process, for both control rooms and flight decks. Preliminary results from exercising one such addition, the real-time cognitive stack display, are presented.

## BACKGROUND

The design challenges facing the power plant control room of the future should seem familiar to observers of trends in cockpit automation. In many ways, the changes contemplated for the advanced control room are analogous to the the "glassification" that has already occurred in both civilian and military cockpits. In the overview that follows, practitioners from the cockpit design world will recognize some familiar issues. These common issues provide a basis for us to learn from to topics of the workshop, and to contribute our observations concerning design principles, including aspects of the design process itself.

Figure 1 shows a notional layout for an advanced control room. In some versions of this concept, there will be no hard-wired controls, although there may well be dedicated locations (either on the large wall display, or on the six– and



**Figure 1.** Conceptual Arrangement of a Computer–Based Control Room

and turbine trips, are already highly automated. An automatic reactor trip leads to a rather different class of scenario where the first steps in an emergency procedure are to carefully inventory that all the automatic functions that are supposed to have occurred on plant trip did in fact occur. In effect, the operators catch up, over several minutes, with what the plant did in a few seconds or less. In the process, of course, they are fulfilling important reasons for having human operators in the first place: (1) to verify automatic actions, and perhaps more importantly given the complexity of the system, to be ready to employ problem–solving as well as procedure–following when appropriate. Thus, a second–level but nonetheless urgent parallel goal is to diagnose the reason for the trip. Typically, following the emergency operating procedures will lead to the steps required to configure the plant to a safe and stable shutdown condition.

In the advanced control room of the future, power plant operators will not change their power settings substantially faster, or otherwise "fly" their plants to new heights, than they do today. Thus the impetus for sybionics in the control room is not the same.

## THE NEED FOR INTERFACE TESTING

Symbionics or no, there are a large number of design issues in human–computer interaction (HCI) that cannot be resolved by appeals to handbooks, yet are too ill–understood to warrant risking large investments, particularly of interface coding effort. Examples include determining the impact of specific interface designs for the HCI–intensive components listed earlier, singly and together. What are the impacts of interface form & function on:

- situation awareness (individual and group)
- operator–paced control vs. event–paced control
- passive monitoring of automatic systems
- response to alarm situations
- control room resource management

Such issues/topics are best adressed by testing proposed interface concepts at the individual, team, and organizational levels. This kind of testing requires integration of three contexts: (1) the task, provided by a scenario that is sufficiently rich to engage (2) the cognitive context, supplied by subject–matter experts who will employ (3) the technological context, supplied by a mix of rapid prototyping and simulation support that are both realistic and complete only to the level needed address the issues at hand.

For example, to what degree does the current prototype interface for computerized procedures offload the operators from the drudgery and error potential involved in the use of paper procedures? Or does the design add secondary tasks that may actually increase mental workload under some circumstances. If so, what are those circumstances?

Testing of the sort just described is very costly in terms of both calendar time and labor. And the supply of subject matter experts is limited. A recent human factors test (different from the one reported here) to investigate sensitivity to control system lags when using soft controls, required several months of preparation to devise a sufficiently–rich scenario, prototype soft control stations, and connect these to an underlying simulation of plant responses. The actual conduct of the test, analysis of data, and report writeup while a relatively small proportion of the effort, took well over a month. This was just one test, and a preliminary one at that, among a planned test series that numbers over a dozen. In most design environments, whether nuclear or aircraft, the number of risk–reduction issues that could benefit from this kind of testing far exceed the resources normally available.

## THE NEED FOR A MODELING TOOL

Thus, if we could computationally model the cognitive consequences of specific details of proposed interface designs, we could both speed and economize the injection of empirically–obtained insights into the tradeoff process early enough to be useful.

The Man–machine Integrated Design and Analysis System (MIDAS) is being developed by NASA[5] to fill precisely this need in the domain of cockpit design: to provide harder ergonomic data to guide the early stages of system development. From a life–cycle point of view, 70–85% of system cost is determined by decisions made in the early stages of development[4].

Conventional systems design processes actively pursue human factors issues late in the process. Typically, much of the information derived from such efforts is categorized as "nice to know", or "operator preference", but not often viewed as evidence for hard requirements. An important reason for this is that many design processes simply cannot afford to view late–arriving human factors data seriously unless the results suggest very severe consequences for operability. Otherwise, the flexibility of the human is relied upon, through selection and training as well as intelligence, to make up for the differences between real–world requirements and those that made it into the system specification.

But if design processes cannot afford to use late–arriving data, it is also difficult in many program cultures to justify a large human factors effort "up front". This is partly due to the reality that the issues may be very dependent on the cognitive characteristics of human tasks that have never been done before, which means that to some extent the knowledge required for the system specification does not exist yet, but rather must be generated as part of the process. This is a qualitatively different approach to design than what works well in cases involving hardware such as bridges and circuit boards. Another difficulty is that the costs of such human factors programs, while perhaps justifiable in life–cycle terms, often seem to loom large at design time. When man–in–the–loop testing is concerned, the costs may include issues identi–

121

of the scenario, one using paper–based procedures, and one using the computerized prototype, were run using the appropriate modules from the MIDAS system (for example, anthropometric and vision modules were skirted).

Results from these comparison runs, showing the effects from the supervisor's point of view, are shown in Figure 3. These results show frozen outputs from MIDAS' new memory stacking display (one of many display windows available from MIDAS, concurrently if desired). The results show a small portion of the 20–minute scenario (about 50 seconds' worth), in which some temporary loading occurred due to the chance simultaneity of several events. These included one operator interrupting the supervisor to correct a prior misunderstanding, arrival of a new alarm signal, difficulty encountered by the second operator in finding an earlier datum requested by the supervisor, and finally, in the case of the paper procedure, a particularly infelicitous formatting decision by the procedure designer.

In this particular incident, many of the memory requirements disappeared because of the computerized procedure's additional information. For example, the computerized procedure inserts current sensor–based plant parameters directly into the text, beside the conditional value called for in the text. The insert is further color coded to indicate (a) whether it represents a violation of the condition called for in the procedure, and (b) the quality of the sensed data based on voting among redundant sensors and paths.

The level of modeling here only shows hypothetical efforts to remember. It is interesting to note that in the actual scenario enacted by human operators, at this same confluence of events the crew proceeded down an erroneous procedure path for over a minute before returning to the correct path. They were unaware at the time that they had taken time for an alignment that was neither required nor effective.

CONCLUSION

The modeled activities, detailed in another report,[7] are perhaps not so significant to this workshop as the glimpse of testing productivity provided by this pilot study. To be sure, productivity from using this display would require prior integration of MIDAS or a similar tool into the design process, so that the incremental cost of accessing or generating relevant





Figure 3. Timelines for Paper–based and Computer–Aided Scenarios.

MIDAS timeline output for the same portions of the MSLB scenario. 3a shows the modeled timelines for the paper–based scenario, while 3b shows the timeline for the computer–aided version of the task. In 3b, one critical incident has disappeared altogether, while the second incident shows far less stacking of remember–activities.

# INTERFACE DESIGN FOR ADAPTIVE AUTOMATION TECHNOLOGIES

P.A. Hancock, S.F. Scallen., and J.A. Duley

Human Factors Research Laboratory
1901 Fourth St. S. E.
University of Minnesota
Minneapolis, MN 55455
USA

## 1. ABSTRACT

Adaptive automation occurs when the control decisions concerning the onset, the offset, and the degree of automation are shared between the human and machine. The purpose of the present experiment was to determine the effect of differing levels of pilot involvement in the the automation invocation procedure. Thirty-two pilots performed a multiple-task battery consisting of tracking, systems monitoring, and fuel management sub-tasks. Monitoring and fuel management sub-tasks were performed manually while the tracking sub-task could be performed manually but was also available for automation. Pilots were assigned to one of four automation invocation conditions. These ranged from complete pilot control to complete system control. Two intermediate conditions involved system-recommended automation and system-invoked automation, where pilots had the opportunity to over-ride the system. Within these groups there were two embedded manipulations. In the first, the shift from manual to automated status was cued by either visual, aural, or a combined visual/aural warning. In the second, the tracking display was reduced in size during automation and relocated either centrally or peripherally. Results indicated that system invoked automation produced less time in manual control, less time to initial automation and an increase in fatigue. Automated tracking display location also affected fuel management performance but this was contingent on the dependent measure used. Results are discussed in terms of consistency of warning signal modalities, the fractionation of performance in multi-task systems, and the increase in workload associated with system invoked automation. A context model for the implementation of adaptive automation is proposed.

## 2. INTRODUCTION

Traditional approaches to automation are founded on task allocation which divides tasks as under the control of either operator or automation [3]. This form of automation can insulate the operator from the system and can take them 'out of the control loop' [12]. This often strips the pilot's role of its meaning and satisfaction. Not surprisingly, pilots are somewhat concerned with automated systems, especially where critical flight decisions are taken from the pilot. McDaniel [5] has commented directly upon the need to retain operator consent for critical decisions such as the launch of weapons. Attaining the perceived benefits of automated systems while maintaining operator authority has consequently become increasingly problematic. Adaptive automation has been proposed as the alternative to static task allocation. In adaptive automation the control of the onset, the offset, and the form of automation is shared between human and the machine [7]. Adaptive function allocation refers to the real-time allocation of function between the human operator and automated subsystems. In the adaptive allocation framework, automation is implemented dynamically, in response to changing task demands placed upon the operator. Adaptive function allocation tries to achieve a major goal of automation, i.e., workload regulation, while avoiding some of the drawbacks associated with traditional static automation. However, adaptive function allocation itself does not identify how the shift from operator control to automated control should occur. There are five major categories which to provide information to implement adaptive allocation. These are triggers derived from: i) pilot performance assessment, using behavioral measures; ii) psychophysiological assessment, using physiological measures; iii) performance modeling, classified broadly as performance models (e.g., signal detection or information-processing

models); iv) critical event logic, using the identification of an unexpected event; or finally; v) a combination system incorporating more than one of these methods [4, 8, 10, 11]. In theory, each of these methods executes some algorithm embedded in the automation system to transfer control. In response to one of the triggers described above, some task or portions of tasks, are automated. While the potential benefits and drawbacks of each method have been identified previously [4, 11], there is an important question looming in the future. That is, the degree of authority that either the operator or the system should have over the invocation of automation. The present experiment evaluates the impact on performance and workload of several of these invocation procedures.

## 3. EXPERIMENTAL METHOD

### 3.1. Experimental Participants, Procedure and Task

Thirty-two experienced pilots (thirty males, two females) volunteered to participate in this study. Their experience ranged from General Aviation aircraft with 100 hours flying time under Visual Flight Rules (VFR) to Commercial Aviation with 16,000 hours flying time under Instrument Flight Rules (IFR). The majority of the pilots fly under Part 135 of the Federal Aviation Regulations. The mean age of the pilots was 32.5 years.

Each consenting pilot completed the Profile of Mood States procedure and engaged in task practice The pilot then completed six, 5-minute trials. At the 4:45 mark of each trial the pilot was asked a series of on-line questions from the Subjective Workload Assessment Technique (SWAT). Following the experimental trials, the participant was asked again to complete the POMS questionnaire. Finally, the pilot completed a subjective questionnaire pertaining to flight experience and the experiment itself. The procedure used the MINSTAR test battery [2], which is illustrated in Figure 1. This presents a multi-task environment in which two-dimensional compensatory tracking, fuel management, and monitoring are presented as individual sub-tasks. In the tracking the goal is to make corrective movements via the flight stick in order to bring the moving cursor in alignment to a fixed target at the center of the display. The difficulty of this task can be manipulated by the experiment. In fuel management, the pilot manually controls the on/off status of the fuel pumps in order to maintain a target level of fuel in the two outer tanks. The

difficulty of this task can be manipulated by initiating failure(s) of the fuel pumps. In monitoring, the pilot is required to reset the lights or gauges whenever they deviate from their normal status.



Figure 1. *The MINSTAR test facility showing the component tasks and their respective location in the fully manual condition.*

### 3.2. Experimental Conditions

The present experiment consisted of a mixed design in which automation invocation procedure was a between-subject factor and warning modality and display re-configuration/location were within-subject factors. The four methods for the invocation of automation for the tracking sub-task were: (1) system-initiated automation (SIA), (2) pilot command by negation (PCN), (3) pilot command by initiation (PCI), and (4) pilot-initiated automation (PIA). In system-initiated automation (SIA), when the pilot exceeded a preset root mean square error (rmse) for a period of 250 milliseconds, the system automated the tracking task. The duration of the automation was 15 seconds after which manual control was returned. Pilot command by negation retained the same performance constraints, however, the system warned the pilot that automation is imminent unless negated. The pilot then may allow the system to automate or negate automation occurrence by depressing a response button within five seconds. For pilot command by initiation (PCI) the system recommended that the pilot automate the tracking sub-task whenever the preset rmse level was exceeded. The pilot may choose to ignore the recommendation and continue performing the task or the pilot may choose to automate by pulling the trigger on the flight stick. In pilot-initiated automation (PIA) pilots used automation at their discretion. An automation

rationale was provided to all pilots which recommended that they use automation whenever they felt performance was deteriorating on the three sub-tasks i.e. when 2500 is not consistently being achieved on the fuel management, not responding quickly and accurately to monitoring deviations, and tracking performance was not constantly maintaining the cursor near the target. All methods with the exception of PIA utilized a constant rmse performance criterion for the tracking sub-task automation initiation.

There were three warning conditions. The auditory warning consisted of a single computer generated "beep" when automation was either invoked (SIA and PIA), imminent unless negated (PCN), or recommended (PCI). A double "beep" was used to signify that automation had been turned off. The visual warning consisted of a message below the tracking display that "automation imminent unless negated" (PCN) or "automation recommended" (PCI). If automation was permitted to occur, the message read "automation on" (all methods). When manual control was returned to the pilot, the message read "automation off" (all methods). The multi-modal warning was a combination of the above two modalities. When automation was either recommended (PCI), warned of its impending occurrence (PCN), and turned on and off (all methods), both the visual and auditory warnings described above were used.

Tracking display re-configuration/location was the final within-subject variable. There were two automated location conditions. When the tracking task was automated the display was reconfigured to a smaller size and then located either centrally (in the middle of the screen) or peripherally (in the bottom right corner of the screen).

### 3.3. Performance Measures

Tracking sub-task performance was assessed through root mean square (RMS) error. Monitoring sub-task performance was given by response time, missed signals and false alarms. Fuel management was scored in terms of deviation from the target value. Four different error derivations were used; absolute error, constant error, variable error and total variability. There were three forms of subjective assessment, these were the profile of mood states (POMS) [6], the Subjective Workload Assessment Technique (SWAT) [9] and an informal debrief questionnaire prepared by the experimenters.

### 4. EXPERIMENTAL RESULTS

#### 4.1. Tracking Performance

Rmse for manual control portions of tracking performance was calculated for each subject and entered as data for the repeated measures ANOVA. Results indicated a main effect for automation invocation ($F[3, 28] = 10.813$, $p<0.01$). Post-hoc tests revealed a significant difference between the system invoked automation procedure and each of the other procedures indicating that the SIA group outperformed all other groups. No main effects were found for the within-subject variables of automation warning modality or display location. The tracking analysis did reveal a significant three-way interaction ($F[6, 56] = 2.462$, $p<0.05$) between all factors.

As the procedures for tracking automation were manipulated as a between-subject variable it was not possible to equate the groups for the total time spent in manual control. The amount of time spent in manual control was calculated for each trial and the data was subjected to a repeated measures ANOVA. The analysis revealed a main effect for invocation procedure ($F[3, 28] = 3.364$, $p<0.05$). Post-hoc tests revealed a significant difference between SIA and PCN groups, indicating that the SIA group spent significantly less time in manual tracking control than the PCN group. The mean time for manual control of tracking by automation invocation procedure is presented in Figure 2.



Figure 2. *Time in Manual Control of Tracking. The difference in time in manual control of tracking is due to the structure of the invocation procedures. SIA automates at the time of violation of the performance criterion, the others allow the pilots to continue manual control at their discretion.*

Rmse data was then standardized according to time spent in manual control. The standardization procedure had two purposes. The first was to evaluate the randomization of subjects to the

125

between-group conditions. The standardized data is a group measure and indicates they are equivalent in terms of tracking ability. The second was to generate a dependent measure of tracking performance which controlled for the between-group differences in the amount of time spent in manual tracking control. Thus, standardized data yield a measure of both tracking ability and tracking performance. Rmse data was standardized by taking individual tracking performance and dividing it by the time spent in manual control on a trial per trial basis. This procedure yielded a measure of tracking performance per unit of time for each trial. This standardized data was then entered into a repeated measures ANOVA. This analysis revealed no significant effects. This indicated that while groups were equated in terms of tracking ability, the effects identified in the rmse analysis are directly due to the different times spent in manual control.

Individual trials were examined to determine the amount of time between the start of the trial and the first episode of automation. This was conceived as an uncontaminated measure of the between-group invocation procedure because the within-subject manipulations (warning and automated display location) are not presented until after automation occurs. Thus, the dependent measure was the time (in seconds) to the first automation episode. Data for each trial was entered into a repeated measures ANOVA with no within-subject factors. The analysis indicated a significant effect for invocation procedure (F[3, 181] = 6.502, $p<0.01$). Post-hoc tests revealed the SIA group was had a shorter elapse time than all other groups, which did not differ between themselves (see Figure 3).



**Figure 3.** *Time to First Episode of Automation. The data indicated that the system invoked automation group (SIA) was exposed to an episode of automation much earlier than the other groups.*

## 4.2. Fuel Management

The level of fuel was averaged for the two goal tanks for each sample of data. Fuel management error was then calculated as absolute error, constant error, variable error, and total variability. In each case the obtained error for each subject, for each condition, was entered into the repeated measures ANOVA. Results for absolute error and total variability revealed no significant main effects or interactions for any of the factors investigated. The ANOVAs for constant error and variable error revealed a main effect for the within-subject variable of automated tracking display location (F[1, 28] = 5.001, $p<0.05$) and (F[1, 28] = 5.788, $p<0.05$), respectively.

## 4.3. Monitoring

Response times, in seconds were obtained for each monitoring deviation and mean response times were calculated for each of the conditions. Mean response times for each subject were entered as data into a repeated measures ANOVA. Results of the ANOVA revealed no significant main effects or interactions. The number of missed monitoring deviations was recorded for each trial, for each pilot. Again a repeated measures ANOVA was completed and no significant main effects or interactions were found. Monitoring data indicated that there were no false alarms. Data also indicated that the missed signal rate was approximately 12.5% per trial but this was artificially high because there were only eight possible deviations per five minute trial.

## 4.3. Subjective Measures

Data for subjective measures of time load, stress level, and mental effort were collected at the 4:45 minute mark of each 5 minute trial for each subject. Data were analyzed via repeated measures ANOVA. The analyses revealed a main effect for the within-subject variable of warning modality (F[2, 56] = 3.924, $p<0.05$). Post-hoc tests, using Post-hoc procedure revealed that the visual warning modality was significantly higher in time load then the multi-modality (combination of visual and auditory). Pre and post trial POMS questionnaires were scored for the six scales according to the instruction manual. A seventh score, reflecting total mood disturbance, was obtained by summing across all scales (scoring Vigor negatively). Pre and post trial scores for the seven scales were subjected to repeated measures

126

ANOVA. Results indicated that Fatigue scores demonstrated a significant main effect for invocation procedure ($F[3, 28] = 3.033, p<0.05$) and a significant procedure by testing interaction ($F[3, 28] = 3.04, p<0.05$). The fatigue data for the main effect indicated that, in general, subjects in the system-invoked automation group reported more fatigue. The fatigue interaction data indicated that system-invoked automation was associated with an increase in fatigue in the post-trial session with all other groups reporting a reduction in fatigue.

## 5. DISCUSSION

If adaptive task allocation is posed as one solution to the problems of contemporary automation, the critical question remains as to how that automation transition is to occur? Tracking performance indicated that the system-invoked automation (SIA) group was significantly lower in their error compared to other groups. However, when tracking was viewed as rmse per unit of time in manual control, there were no significant differences between invocation procedures. This suggests that it is not capability but strategy that is changed by invocation differences. This was confirmed by examining 'the time to first automation.' This is a pure measure since within-subject factors warning and automated display relocation were not presented until after the first automation. Data indicated that the system-invoked automation (SIA) group were transitioned to automation much more quickly than the others. Thus, a fundamental effect associated with invocation procedures and the major finding here is that system-invoked automation is associated with less time to initial automation.

Results from the fuel management task address the human-machine interface as it relates to the display of automation status. The specific conclusion, however, is related to the question, "What effects occur, in terms of performance bias, for component task fuel management, when primary task tracking is automated to display position either centrally or peripherally"? The answer, in terms of bias, is that the fuel management performance favors the peripheral automated tracking display. "What effects occur, in terms of performance variability, for component task fuel management, when primary task tracking is automated to display position either centrally or peripherally"? The answer, in terms of variability, is that fuel management performance

favors a centrally located automated tracking display. Which conclusion is more valuable? The answer depends on how critical the respective types of errors are in the real-world. If one seeks to control individual variability then a central automated display favors better performance on other component sub-tasks. If, however, one seeks to control bias error across a group of users then a peripheral automated display factors better performance on other opponent sub-tasks. These conclusions are reflective of individual differences, primarily because variable error is computed with the individual performer's mean. Finally, a practical application of the fuel management data is that fuel management performance was not specifically related to the invocation procedures under examination or, more generally, to the concept of a continuum of authority (from fully manual control to fully automated control) as it applies to the invocation of automation.

With respect to warning modalities, it was proposed to contrast modes of sensory assimilation. The hypothesis being that multi-modal warnings would be most effective. In constructing an ecologically valid display it was not possible to provide psychophysical equivalence between visual and auditory cues. However, results are informative for the practical realm. It appears that simplicity should guide design in invocation advice in terms of interface development. It was clear that visual, alpha-numeric communications were the least preferred conditions by pilots and that some degree warning 'horse-race' was imposed in an information-processing dynamic. Our recommendation for simplicity also argues for a careful evaluation of other systems placed in the cockpit which might signal the onset or offset of critical functions.

### 5.1. A Context-Based Model

After extensive investigation of the question of adaptive allocation, it has become clearer that there is no single right answer. The answer always depends upon the context of performance. Consequently, we view the adaptation process of automation in terms of envelopes of protection. What drives the system are the needs to achieve mission goals. What bounds the system are the constraints of the operator, the flight platform, and the flight conditions they encounter, more generally the environmental constraints.

**Figure 4.** *A context model for the implementation of adaptive automation systems. The contours represent zones of operation which differ in the degree of operator control. The path of optimal progress is human-centered but interaction with environmental constraints may force the operator to relinquish degrees of control. Note that there is no human-centered contour on the very right. It is unlikely that multiple exposure to environmental constraints would allow the operator to maintain total control. At such junctures, consultation or emergency (life-threat) strategies may be initiated.*

The following model is a descriptive but dynamic view of adaptive allocation (see Figure 4).

Initially, environmental constraints are not severe and the range of operational latitude allows each of three allocation strategies to be engaged. As environmental constraints are loosened (as for example in straight and level flight), operational latitude increases and the ability to use human-centered strategies is enhanced. As environmental constraints are tightened (as for example in pilot incapacitation, engine failure, or nap-of-the-earth flight) possible strategies are reduced. The outer region of operational latitude is represented as a life and system-threat envelope. Conditions which threaten to violate these bounds (e.g., flight into terrain), are dealt with by the system-level of automation initiation. This is particularly the case in emergency conditions such as the incapacitation of either pilot or some vital aircraft component. Embedded in this operational envelope is a region of consultation. This is typified by the interactive strategies of command by negation and advisories. Finally, the inner region is human-centered in which all automation decisions are initiated by the pilot. We would like to maintain the pilot in this region of optimal control, however, many environmental contingencies vary which act to perturb the path of progress. These perturbations are communicated in terms of tasks to be performed and hence re-configuration of the multi-task matrix

is periodically needed. The present descriptive model has the advantage of combing all forms of automation initiation strategy dependent upon context and preserving safety of operation in the outer envelope.

REFERENCES

1. Fitts, P.M. (1951) (Ed.). *Human engineering for an effective air navigation and traffic control system.* Washington, D.C., National Research Council.

2. Hancock, P.A., Scallen, S.F., & Duley, J.A. (1993). *Initiation and cessation of automation: location versus configuration change* (Technical Report). Naval Air Warfare Center, Warminster, PA.

3. Hancock, P.A., Scallen, S.F., & Duley, J.A. (1994). *Pilot performance and preference for cycles of automation in adaptive function allocation* (Technical Report). Naval Air Warfare Center, Warminster, PA.

4. Hilburn, B., Molloy R., Wong, D., & Parasuraman, R. (1993). Operator versus computer control of adaptive automation. *In the adaptive function allocation for intelligent cockpits (AFAIC) program:*

*Interim research and guidelines for the appellation of adaptive automation.* (Technical Report NAWCADWAR-93931-60). Naval Air Warfare Center, Warminster, PA.

5.    McDaniel, J. (1988). Rules for fighter cockpit automation. *Paper presented at the IEEE National Aerospace and Electronics Conference.* New York: IEEE, 831-838.

6.    McNair, D.M., Lorr, M., & Droppleman, L.F. (1971). *Profile of mood states manual.* San Diego, CA: Educational and Industrial Testing Services.

7.    Morrison, J.G., Cohen, D., & Gluckman, J.P. (1993). Prospective principles and guidelines for the design of adaptively automated crewstations. In, *The adaptive function allocation for intelligent cockpits (AFAIC) program: Interim research and guidelines for the appellation of adaptive automation.* (Technical Report NAWCADWAR-93931-60). Naval Air Warfare Center, Warminster, PA.

8.    Parasuraman, R., Bahri, T., Deaton, J.E., Morrison, J.G., & Barnes, M. (1990). *Theory and design of adaptive automation in aviation systems. (Technical Report).* Warminster, PA: Naval Air Development Center .

9.    Reid, G.B., & Nygren, T.E. (1988). The subjective workload assessment technique: A scaling procedure for measuring mental workload. In P.A. Hancock & N. Meshkati (Eds.). *Human mental workload.* (pp 185-218). Amsterdam: North Holland.

10.    Rouse, W.B. (1988). Adaptive aiding for human/computer control. *Human Factors, 30,* 431-443.

11.    Scallen, S.F., Duley, J.A., & Hancock, P.A. (1994). Pilot performance and preference for cycles of automation in adaptive function allocation. *In Proceedings of the First Automation Technology and Human Performance Conference,* Washington D.C., April.

12.    Wiener, E.L. (1988). Cockpit automation. In E.L. Wiener & D.C. Nagel (Eds.). *Human factors in aviation* (pp 433-461). San Diego: Academic Press, Inc.

# Displays and Controls for the Pilot/Electronic Crewmember Team

Richard Lynch
Christopher J. Arbak
Barbara J. Barnett
John L. Olson
McDonnell Douglas Aerospace-East
New Aircraft and Missile Products Division
P.O. Box 516, M/C 0642203, St. Louis, Missouri 63166

## 1. SUMMARY

The quality of communications between a pilot and an electronic crewmember (EC) will help determine whether the EC becomes an accepted and trusted member of the pilot's onboard team. This paper presents McDonnell Douglas Aerospace (MDA) efforts relating to the design of a communications interface between a pilot and an aircraft with an EC.

Our objective for the pilot/EC interface effort was to apply what is known about cognitive psychology, human factors, and design techniques to the design of an intuitive information interface for the pilot of an EC-equipped multi-mission fighter. After researching and interpreting these topics, we identified principles for interface design most applicable to the user population. Next, we used those principles to develop several pilot/EC display formats. We installed these formats into a simulator and used an emulated EC to conduct man-in-the-loop evaluation runs. We examined the results of the evaluation and documented the lessons learned. These lessons can be applied to future interface design efforts, not only for pilot/EC teams but for crewstation configurations in general.

## 2. INTRODUCTION

The Advanced Crew Systems group at McDonnell Douglas Aerospace has conducted research since the early 1980's in the area of optimizing cockpit displays and controls to help increase aircrew situation awareness. In the last few years this research has expanded to include interfaces between a pilot and a notional avionics subsystem that can assume many tasks of a second crewmember. Our goal of providing intuitive displays and controls for the pilot armed with such a system was based on the philosophy that effective cockpit communications between the pilot and his electronic crewmember is one key to increasing acceptance and trust in a technology that is often stereotyped as intrusive or unnecessary. The removal of these labels is essential to building trust in the system, and is possible if the appropriate roles are established for each member of the team. The premise we adopted for establishing roles is that the pilot is always in charge, will always be held accountable for all decisions, and therefore must shoulder the greatest amount of decision-making responsibility. While this should seem like an obvious fact, there is always the danger of ignoring it and designing a system which does not work for the pilot but makes the pilot work for it. With the correct focus, the EC is simply one of many tools the pilot can exploit to successfully accomplish his mission.

MDA has developed and evaluated displays and controls for a pilot/electronic crewmember team based on these concepts. Our assumptions of the electronic crewmember's capabilities, the principles of design we used to develop the interface, and the evaluation environment will be described in this paper. Lessons learned from multiple simulation runs with experienced fighter pilots will be related.

## 3. ELECTRONIC CREWMEMBER FUNCTIONALITY

To adequately evaluate pilot/EC interface designs, we estimated what capabilities a future EC-like system would realistically be expected to have. To establish the foundation of our interface designs, a number of initial guidelines for EC capability were developed around our philosophy of the EC as an intelligent subordinate to the pilot. Specifically, the EC:

- could not act on its own;
- could make recommendations;
- could take action based on pilot direction;
- could take action based on interpreting pilot intent;
- could fly the aircraft tactically on autopilot;
- could deal with ambiguities in human speech in the context of the mission;
- could diagnose malfunctions, identify miscommunications, and determine the correct response.

Our goal was for the EC to provide consistently correct information and aid the pilot's decision making by helping to manage workload, reduce confusion, and simplify tasks. The first of these traits requires the EC to be virtually error-free. With an actual electronic crewmember, confidence leading to trust can only be achieved after continued success at providing assistance that is both logical and expected by the human, and is correct all the time. Depending on the situation one error by the system may strongly bias pilots against using it at all, since pilots are much less forgiving of systems errors in flight environments filled with critical mission phases and timing constraints.

The electronic crewmember also must not contribute to pilot workload, but should help him manage what he already has. This may not reduce workload, but high workload is not the main problem because pilots generally prefer to stay involved in operating their own aircraft. This helps keep their concentration level up, allowing them to "stay ahead of the jet". The real problem is task saturation, where low situation awareness or strict time limits may cause the pilot to get behind and not know exactly how to catch up. Task saturation tends to follow a snowball effect that at best causes the pilot to fall even further behind, and at worst may cause him to "freeze up". To combat this problem, pilots strive to gain the knowledge and experience needed to know what is going on and what to do next in the most probable extensions of the current situation. An electronic crewmember that can consistently demonstrate the ability to help the pilot stay ahead of his jet while ensuring he stays in the loop as mission manager will be readily accepted and trusted.

## 4. INTERFACE DESIGN PRINCIPLES

Once the capabilities of our electronic crewmember were agreed upon, the design team began the process of developing the display and control designs we would evaluate in simulation. Our design team consisted of both experienced fighter pilots and human factors engineers, allowing us to keep an operational perspective while ensuring that we considered known cognitive psychology and human factors design techniques. In the initial stages of the process, these techniques were developed into the principles for design we would use on the pilot/EC interface. These principles were based on cognitive characteristics unique to the fighter pilot population in order to make the designs user-friendly and "intuitive". In the context of fighter displays and controls, an interface that is intuitive is easy for aircrews to use and quick to learn, where the required actions are obvious, implications of those actions are clear, and user expectations of performance are met.

Similar principles were grouped together so that three main categories emerged, dealing with spatial representation, perception and cognitive tasks, and itegration of data. Within each group were several general principles we found most useful for the pilot/EC interface designs:

### 4.1 Spatial Representation
- Represent information consistent with the spatial environment. In the cockpit, the environment is from the pilot's point of view inside-out and is scaled.
- Display current location in the mission as an indication of position within the sequence of tasks. Pilots tend to separate missions into distinct segments, each with special groups of functions to be performed.
- Represent the complexity of the environment, including all data relevant to the pilot's decision-making process. Attempts to simplify the interface by arbitrarily limiting the data which reaches the pilot runs the risk of lowering situation awareness.

### 4.2 Perception and Cognitive Tasks
- Present information graphically for holistic processing. Pilots can assimilate the "big picture" more rapidly through the use of graphics in a crew-centered design.
- Use alphanumeric information to facilitate detailed, analytic processing. Discrete use of textual data will enhance the quality of the graphical information.
- Use comparative, not absolute, judgments to aid rapid decision-making in time critical missions.
- Use specialized displays for specialized tasks.
- Use color as a redundant code to facilitate rapid processing.
- Display important information in consistent locations, lessening the time required to look for the data.
- Automate tedious, repetitive tasks.

### 4.3 Integration of Data
- Combine related data into a representation that groups the information into meaningful units. Data scattered about a display without regard to how it relates to other pieces of information increases workload and crosscheck time unnecessarily.

- Reduce visual clutter and minimize the number of items to be searched in a display. Gestalt principles can be used to ensure all relevant data is provided in the most efficient and uncluttered way.
- Highlight important information or information contrary to the norm, and call attention to neglected items.
- Provide feedback for control inputs.

## 5. MANNED SIMULATION DEVELOPMENT

Before we applied these principles to actual display and control design, a mission was chosen for the simulation evaluation with events such as a low level ingress to an air-to-ground attack, enroute threat engagements, systems malfunctions, and retargetting tasks. Due to time constraints we chose not to model an entire mission with every contingency, but concentrate on segments most interesting to a study of the pilot/EC team. Our approach was to design the mission and script the probable dialogue between the pilot and EC, referencing our assumed EC capabilities, in order to determine which specific displays to focus our design effort toward. By scripting the mission we were able to bound the evaluation within reasonable limits and avoid problems with pilots venturing off into areas we hadn't designed an EC interface for.

We based our electronic crewmember emulation on the trade-off between the cost of developing a real electronic crewmember and the need to simulate a realistic, useful, and trustworthy system. We found that the best way to meet all these requirements was to use experienced fighter pilots playing the role of the electronic crewmember, which gave the system instant credibility since all pilots who participated in the simulation were confident the "electronic crewmember" would meet their expectations of performance. This has been necessary in order to focus on designing the interface rather than worrying about the operations of the EC, since our goal has been to develop an easy to use and understand interface allowing two-way communication with the electronic crewmember.

To gain insight into the intuitiveness of our display or control designs, we evaluated them in the proper environment with a cross-section of experienced users. For a pilot/electronic crewmember team, this was a manned simulation of a high workload mission with many opportunities for the team members to interact. System malfunctions, threat attacks, mission replanning, and coordination with higher authorities were all elements that we included to stimulate the communication between the pilot and the electronic crewmember. Even normal mission events such as low level flight, sensor operations, target area attack, and egress were areas where the communication seemed to naturally occur as the pilot worked to gain the knowledge to build his situation awareness and reduce task saturation.

## 6.0 DISPLAY AND CONTROL DESIGN

Once the A/G attack mission segments were selected and scripted, the design principles were used to guide the development of several display formats which the EC would use to communicate information and recommendations to the pilot. The principles were also used to help determine the optimum controls for the pilot to command the EC.

### 6.1 Interface configuration
Several of the following cockpit enabling technologies were considered for our pilot/EC interface designs, based on their

ability to improve pilots' situation awareness and reduce task saturation:

- Large head down displays
- Helmet mounted displays
- Full color capability
- Display windowing
- Graphical display capability
- Voice control
- Touch control
- Hands-On Throttle and Stick (HOTAS) control

## 6.2 "Cockpit 2001"

The flexibility and large display space of the McDonnell Douglas Cockpit 2001 concept provided important benefits to our pilot/EC interface. As shown in the following figure, Cockpit 2001 features two 10 inch by 10 inch flat panel displays mounted side by side on the main instrument panel. This results in more than twice the programmable display space available in current fighters such as the F-15E or F/A-18. With windowing capability the pilot can select several display formats simultaneously if desired, or he can select a large, global "panoramic" display of the mission environment. Hands-on "macro mode" control allows him to select format configurations tailored to each mission phase. For our primary EC control interface we simulated a



**Figure 1. Cockpit 2001 Display Configuration**

voice control system with natural speech capability to provide a degree of familiarity and personality to the team. We also simulated HOTAS, touch control, and standard pushbutton/switches as alternate methods of allowing the pilot to communicate with the EC.

## 6.3 Pilot/Electronic Crewmember Interface Designs

With Cockpit 2001 as the basis, we began to apply the design principles both to a general EC interface structure and to several specialized formats for communicating the appropriate information on scripted mission events. To satisfy the basic principle of consistent location of important information, we dedicated a central message location on the Up Front Control for initial EC communications and prompting. We also color-coded formats and

recommendations generated by the EC, to differentiate from normal systems data. This was intended to reduce the amount of time the pilot needed to search his displays to communicate with the EC. Magenta (light purple) was selected, since it is used sparingly on current displays, and would not cause conflicts by being devoted to the EC. Figure 2 shows the Up Front Control with prompting "RETARTETING . . .READY" displayed inside the magenta EC message window. A redundant voice message accompanied all UFC messages, helping to let the pilot know there was data he needed to look at.



**Figure 2. Up Front Control Message Window**

If an EC message appeared on the UFC that had accompanying detailed data for the pilot to view, the pilot could select an EC window on one of his large displays to examine the information. These windows also featured a magenta border, and any EC-unique data or recommendations were also displayed in magenta.

The next few paragraphs summarize a typical mission script and illustrate the EC interface formats the pilot used to deal with the various mission events.

During the initial low level ingress segment, the EC detected a system malfunction and alerted the pilot through a warning on the Up Front Control EC message line. While the EC diagnosed the problem and determined the best course of action, the pilot was presented with status information and was asked permission prior to any change in system configuration. By applying the principles dealing with holistic graphics, specialized displays, and highlighting data out of the norm the three formats in the following figure were generated. The display in the lower left was the first offered by the EC and shows a left generator failure. The EC asked permission to recycle generator power in an attempt to restore normal operation. The next format in the lower right depicts the diagnosis of a fuel leak which also imposes limits on maneuvering and afterburner usage. In the third, upper left format the EC offered additional details about the problem, including the specific location of the fuel leak, rate of fuel loss, and a recommendation to set idle power to reduce the loss rate. The pilot selected each of these formats whenever he was ready for the information, and the EC kept them in sequence to guide the pilot through the problem.

132

**Figure 3. System Malfunction Formats**

In a later segment of the mission. the EC received inputs from the simulated onboard sensor integration system that enemy aircraft had been detected and were on an intercept course to the friendly strike package. The lower left format in Figure 4 relates the detected enemy formation, speed, and altitude, and the EC resolution of aircraft type and probable weapons carriage. In the mission script, the mission commander in a different aircraft decided that "blue" flight, which was ahead of ownship, should engage the enemy while the rest of the package continued the mission, so the EC related this to the pilot through the UFC message window. The mission commander also wanted to know if ownship could retarget to "blue" flight's target, and the EC displayed the information necessary for the pilot to make this decision in two additional formats. In the lower right the EC presented the target environment graphically, and highlighted blue flight's target to the pilot. In the upper left, the EC showed the pilot that fuel, weapons, and time on target criteria were OK for retargeting, and asked the pilot to confirm or refuse the retargeting request from the mission



**Figure 4. Retargeting Formats**

commander. In the figure, the pilot has accepted the new assignment, so the "Retargeting OK" bar has changed to solid green from a magenta outline.

While approaching the final run-in to the target area, the EC calculated the two best attack axes taking into account threat positions, terrain, sun angle, and other tactical considerations. In the lower left format of Figure 5 the pilot has already selected the zone which allows three targets to be attacked in rapid sequence using a multiple target attack algorithm. The integration of graphic and alphanumeric displays of the attack plan allowed the pilot to quickly assimilate the required information (number of targets, inbound attack heading boundaries, orientation to the target area) and make a timely decision. The lower right format shows the pilot how his final attack run-in appears so that he could orient himself early, and know what to look for to positively identify the target. The format also shows



**Figure 5. Multi-Target Attack Formats**

through color coding when his heading was within the proper limits for the algorithm to release against all three targets. The upper right format is a combination of SAR imagery of the target area, overlaid with the three targets recommended by the EC and selected by the pilot for attack.

**7. LESSONS LEARNED**
We determined three roles in the pilot/electronic crewmember team from observations of the subjects: mission manager, systems manager, and flight path manager. Every pilot we tested gave the flight path manager role to his electronic counterpart as soon as workload began to increase and decisions had to be made. Since the pilots trusted the aircraft to be flown tactically, they turned their attention to mission management tasks. The EC was even allowed to fly the attack in the target area in order to release weapons against multiple targets. The systems manager role tended to be split, with the electronic crewmember diagnosing problems and recommending fixes, while the pilot reserved the right to give permission to actually change system configuration if required. Thus, we learned the systems manager and flight path manager's roles can be accomplished by the electronic crewmember while still being dependent on decisions from the human mission manager.

133

Five levels of pilot/EC interaction may be desired, tailorable by mission phase, systems involved, and threat expected. We have found pilots differ as to how much they want the EC to be able to do on its own. The five levels are:

- Fully automatic EC with no feedback to the pilot
- Fully automatic with feedback on status
- Semi-automatic EC (tells pilot what will happen and asks permission to proceed)
- Manual (provides information only; pilot is expected to perform all tasks with no assistance)
- No EC involvement

The pilots commented that the interface must facilitate rapid review and acceptance or tailoring of these options, to be useable in time-critical situations.

More head up cues are desired. Pilots typically don't stare down into the cockpit, and don't want to have to monitor and control an EC that way. Pilots' desires to remain head up during a low level tactical air-to-ground mission plays a signficant role in CVI design, including EC interfaces. Voice control, helmet mounted or HUD cues, or an up front control, all allow varying amounts of the desired head up/head forward operation.

Voice was the control of choice, due to a combination of limited training, ease of use, and flexibility of the emulated EC. In a typical simulation mission lasting 10 minutes, 600 exchanges between the human pilot and the electronic crewmember took place, of which 500 were through the emulated voice control system. At times during the mission the EC had relatively large amounts of data to communicate to the crew, so that certain display formats were crowded. The pilots typically spent little time trying to interpret these cluttered displays, and instead fell back on the voice control capability to ask for the most important data and recommendations. The pilots appeared to function more easily with multiple simple formats showing specific pieces of data, easily requested and exited using voice control.

In a fast-paced, workload intensive scenario, prompting from the EC was very helpful in keeping the pilot thinking ahead, thus helping to prevent task saturation. The key element here is helping the pilot make the mission decisions in a timely manner. To reduce crosscheck requirements, an ideal alerting system would provide a maximum of immediately useful information in as little time or space as possible.

Options or alternatives were generally understood when offered together for comparison. If the options or their method of display was complex, they were often overlooked, not noticed or misunderstood unless studied extensively which the pilots did not want to do because of the danger of

getting behind. A pilot that trusts his EC will seldom examine the alternatives to the primary recommendation, thus reducing time spent making decisions. Lack of trust will force examination of the options, decreasing the benefits of workload management afforded by an EC. Lack of trust can lead to a counterproductive and potentially dangerous situation as the crew runs out of time to do all the tasks, causing task saturation.

Pilots often had multiple EC formats displayed at once on the Cockpit 2001 displays. They would tell the EC which formats they wanted to keep and which could be removed, and in this way keep situation awareness high for each mission phase. One important aspect of the pilot/EC interface is the desire to have an easy way to remove specific EC data when it has been assimilated and is no longer needed. We found this to be a non-trivial task, especially when combined with the concept of multiple, simple formats displayed at once. The selection of an EC format was easier, since the formats were queued and selected in order as the pilot was ready to receive them. The problem with removing formats is that the pilot needs an easy way to tell the EC which particular format he doesn't want anymore. Any CVI design for pilot/EC must make this a priority.

Human emulation of an EC is an acceptable method to study the interface, as long as the people playing the EC are a part of the same user population as the test subjects, so that there is credibility and expectations of performance are generally met. In this way, the EC will be able to interpret ambiguous pilot communications and make its response unambiguous, and also realize when misunderstandings have occurred and correct them.

## 8. CONCLUSION

A human/electronic crewmember team can be trusted to successfully accomplish the mission, assuming the appropriate roles are assigned to each team member, the EC is dependable, and a user-friendly and intuitive interface between the two exists to allow smooth communications in time-critical, workload-intensive situations. Graphical displays, prompting, highlighting, interactive voice control, and specialized formats have all been shown to be very effective features of a pilot/EC interface. While futuristic, large display configurations enhanced the usability of the EC interface formats in our evaluations, the same design principles identified in our work can be used to develop formats intuitive to pilots of current day aircraft. Designers must also remain aware of our findings of the pilot's desire to be the overall mission manager while allowing the EC to be the systems and flight path manager. Although based on an emulation of an advanced, "intelligent" EC, these results should nonetheless provide guidance to future design efforts.

# Standards for Trustworthy Human-Electronic Teamwork

B.M. Sherwood Jones
BAeSEMA
1 Atlantic Quay
Broomielaw
Glasgow
SCOTLAND
G2 8JE

## INTRODUCTION

This paper examines the need for standards for trustworthy Human-Electronic Crew (H/E Crew) teamwork, and the feasibility of setting and implementing such standards; what do we need to do before we can deliver H/E Crew systems with confidence? The paper is based on the belief that the H/E Crew is "special" - the whole is more than the sum of the parts - and that additions to existing standards and engineering methods will be needed to deliver trustworthy H/E Crew teamwork. The author's experience is that there is now significant customer pull for intelligent support systems (under a variety of names) and an expectation that such systems can be delivered at affordable risk and cost.

## WHAT IS SPECIAL ABOUT THE H/E CREW?

Human-electronic systems operating at the cognitive level are different to systems that operate at the physical and physiological level. Hollnagel and Woods (1) proposed that such systems are something "special" over ten years ago.

A cognitive system (Woods, 2) is goal directed; it uses knowledge about itself and its environment to monitor, plan, and modify its actions in the pursuit of goals; it is both data-driven and concept-driven. More significant than any changes in technology, what is being proposed is a step up in the level of description to a *knowledge system* (Newell, 3), rather than a program-level system.

It is proposed that progress in developing H/E teamwork can be compared with progress in understanding flying qualities. Vincenti (4) has described the learning process (over a quarter of a century) by which an ill-defined problem with a large subjective human element was translated into an objective well-defined problem for the designer. One of the striking parallels was the move from an open-system viewpoint - which saw the pilot as an external agent - to a closed-loop system which saw the pilot and aircraft as a single system. Although there is a community of engineers and researchers who assume the closed-loop total-system approach, it is in the minority. Further, within that community, the potential problems raised by Woods (2) eight years ago are perhaps better understood, but they are not necessarily nearer resolution.

## THE NEED FOR STANDARDS

If we are to conduct engineering at the level of knowledge systems, then we need to develop appropriate standards just as a matter of good practice ("everybody needs standards"). However, knowledge level systems have their own requirements, and these pose specific demands on the engineering process. This is not equivalent to any specific technology such as "KBS", and to some extent, it could be argued that the failure of many current computer systems is due to a failure to recognize the knowledge level (the long history of failure in 'decision aids' is a case in point). Perhaps the 'H/E team' is more widespread than we think.

The specific requirements of the H/E team include the following; dialogue about shared goals and values, handling of uncertain and conflicting data, shared data interpretation and decision making, shared error recovery (to include 'mistakes' as well as data entry 'slips'), the need to interpret and accept machine decisions, the ability to distinguish plausible wrong answers, and joint adaptation or evolution as part of organizational 'self design'. A key ingredient of interest to this audience is real-time operation, e.g. on-line plan repair rather than off-line planning.

## PRESENT STATE OF H/E CREW TECHNOLOGIES

There have been major advances in the (necessary but not sufficient) technologies that support H/E Crew teamwork. This section gives a thumb nail sketch of where relevant advances have been made, and where we still have a way to go before building H/E teams can be considered a regulated process.

### Human Factors, HCI

Human Factors in general, and UK Military HF in particular, has become much more regulated and standardised. The integration of HF with the rest of systems engineering is less well-defined. The specific area of Human computer Interaction (HCI) has become much more standardised (e.g. 5). Some parts of ISO 9241 (5) have an oblique relevance to the H/E Crew, but no particular consideration is given to any form of intelligent support.

### Software Engineering, Cognitive Science

The dominant paradigms in both cognitive science and software engineering are particularly unhelpful to the development of H/E teams:

Software engineering still takes little account of the user as part of a total system, and most work on safety-critical systems (where one might expect trustworthiness) is focussed entirely on the machine element, and to some extent the role of the operator as fall guy for design inadequacies is becoming formally recognized.

If the Newell approach to cognitive science (3) is taken as representing the dominant paradigm, then the H/E Crew would appear to be impossible; the social band operates on a timescale of days to months. "There is no way for a social group to act as a single body of knowledge. This restriction applies to goals as well. Whatever advantages may be found in reconciling opposed or discordant goals, the limited means of communicating about goals and why they exist is sufficient to keep the organization from ever arriving at a single common basic goal structure." (Newell, 3). This is greatly at variance with the findings of the High Reliability Organization programme (6, 7) and most people who have studied well-formed teams. The Knowledge system is portrayed in just the open-system manner that early work on flying qualities portrayed aircraft stability. The cognitive science paradigm also gives no separate credence to affect, emotion or the priorities of survival (motivated by the autonomic nervous system in the human element), and hence to the 'prime directives' or defaults proposed by Reising and Emmerson (8) (based on Asimov's laws of robotics).

### Human teamwork, decision making

The human side of the team has been more promising; The last few years has seen the continuing application of Cockpit Resource Management and an increased research interest in understanding command teams and how they work (e.g. Sherwood Jones 9, 10, and other papers in this symposium). The key message at this stage is to stress how much more is required for effective teamwork than just procedural information processing. The other major development on the human side has been the huge change in our understanding of human decision making brought about by the Naturalistic Decision Making school (the work of Gary Klein and others).

### System Engineering, Safety Management

There are two developments of relevance as regards the overall approach to systems engineering. (Important knowledge engineering developments, such as the increasing practicality of KADS are not relevant to this discussion, since they are largely at the program-level (3) rather than the knowledge-level).

The first topic is a non-development; there is still very little progress as regards designing systems for quality of decision making (rather than speed of response or 'workload').

The second development is much more helpful. This is the increased application of hazard-driven design approaches under a safety management framework (Def Stan 00-56 (11), Carr et al, 12). Coupled with this is the increased recognition being paid to organizational issues in system safety.

Consideration of the Individual, the Job and the Organization is now a necessary part of producing a safety case in the UK, and the Management of Health and Safety at Work Regulations mandate such a safety case where there is a hazard to users or other people affected by the system. It is considered that the best way to make progress at the moment is to develop H/E Crew-specific hazard checklists for use in support of a safety case (e.g. to supplement the Hazard Identification Checklist and assist Operating and Support Hazard Analysis under 11). One of the prerequisites for a High Reliability Organization (6) is redundancy in decision making capacity, and the Electronic Crewmember may be the way of achieving this redundancy.

### STANDARDS NEEDED

The aim of a design *process* standard such as the one proposed is to provide developers with a systematic process for identifying potential teamwork hazards and eliminating them through design. The standard would constitute a statement of "best practice." If developers can demonstrate that they have followed the best available design methods, it will follow that they have given appropriate attention to safety.

The first step would be to eliminate from further consideration, those systems that do not operate at the knowledge level (and which therefore do not raise teamwork or trust issues). This involves examining the activity to be undertaken by the H/E team, (probably being currently performed by humans). If questions such as the following are answered "yes", then the system needs to be considered at the knowledge level.

Are operational goals represented or implied? Would the behaviour or output be expected to vary if operational priorities change?

Are uncertain or conflicting data reconciled, and action taken? Is additional data sought to resolve data inadequacies?

Are reports or explanations given to account for actions taken?

Are experienced users/teams better than ones straight out of training? Has the method of operating changed in the last year? Do operating practices vary from place to place?

It may also be useful to add the traditional Expert Systems questions about whether the problem is normally considered to be a demanding human task, but this is unclear at the time of writing.

### THREATS TO TRUSTWORTHY TEAMWORK

Having identified that the human-electronic system needs to be considered at the knowledge level, then the design and test of the teamwork needs to be considered against a number of threats to trustworthiness. The tentative presentation of threats below is based on the Competing Values Approach to teamwork of John Rohrbaugh (13, 14) which is based on sound theoretical foundations and wide-ranging empirical refinement and validation. The decision making process (rather than outcome) is

assessed. Effective teamwork needs satisfactory performance when judged from four perspectives (reflecting the different Competing Values). The perspectives are summarised below, with a few of the H/E Crew specific points added.

The **rational perspective** (or rational goal model) is concerned with on-line information processing under conditions of high volumes of data. Here we need to check a) Is the process sufficiently goal-centred? Are goals clear and understood? Is there a logical framework for evaluating courses of action? b) Is the process efficient enough to cope with the volumes of data in the time available? Is the cost of arriving at a course of action too high? It is here that the inadequacies of our understanding of human decision making show up; we do not have an ordered set of decision models and when to use them, or how to support them. Some progress on how to handle specific types of human error has been made (Silverman 15). Protocols for eliciting probabilities have reached a level of maturity whereby the major threats can be countered (16). The most frequent failure (in the author's experience) is the use of point values rather than distributions. The technical performance must be predictable enough to be trust*worthy*.

The **empirical perspective** (or internal process model) is concerned with information management, or planning and control processes, including the dynamic allocation of function to achieve a sense of order and smooth functioning. Here we need to check a) Does the team have adequate access to the evidence it needs for decisions, is information distribution timely, are there adequate information management procedures? b) Is the decision process traceable and accountable, can it be reported, justified and explained? How an Electronic Crewmember can be held accountable is a long-standing concern. A software 'functional' engineering approach can easily miss user information needs and thereby prevent effective action (17). A common design failing is consider a system as an isolated platform, rather than as part of an operational setting. Providing explanations that meet the various needs of different levels of the command chain may turn out to be a major design driver (17). The H/E team must be dependable enough for the development of trust.

The **consensual perspective** (or human relations model) is concerned with coordination processes to achieve a skilled, cohesive team with high morale. Here we need to check a) Is the decision process sufficiently participatory, do all members of the team get their say, are their interests represented? b) Is consensus achieved, are outcomes supported? Good information processing does not necessarily lead to user acceptance. Letting both human and electronic team members to have their say and be happy with the result places great demands on the interface. The problems of supervisory control are well established, but there is still a 'functional' tendency to create monitoring jobs. The consensus-building must be good enough to develop trust*edness*.

The **political perspective** (or open systems model) is concerned with the management processes of flexible adaptation to a changing environment, including drawing on additional resources when needed. Here we need to check a) Is the process sufficiently adaptable, can procedures be altered, handle changes in the situation (i.e. provide requisite variety) b) Does the process take account of the interests and concerns of external groups, follow ROE's, fit in with command priorities? Are proposed actions considered from the enemy, ally, neutral viewpoints? Is the survival of comrades properly considered and weighed? This is where the all-pervading defaults and prime directives (8) come in. These are not 'rational', 'cognitive' goals, but are concerned with fight and flight - they are affective (18). Respect for fiduciary obligations is necessary for trust*worthiness*.

## PRESENT LEVEL OF UNDERSTANDING

Vincenti (4) described a number of interacting elements associated with the improvement in understanding flying qualities. These elements seem an appropriate framework in which to describe the extent to which H/E teamwork is a fully mature engineering process. This section lists the key elements (with very slight adaptation) together with the author's impression of the maturity.

1) Familiarization with H/E team and recognition of the problem; there is a body of literature on the potential problem (e.g. Woods, 2), and the beginnings of a literature on experience of the problem as experienced on demonstrator or experimental systems (Sloane, 19, Sherwood Jones et al, 17) but mostly on systems that were *not* implementations of the teamwork concept. The body of this experience is likely to grow rapidly over the next few years.

2) Identification of basic variables and derivation of analytical concepts and criteria. We are still at the early stages of this. The parameters of human teamwork are beginning to be understood, the parameters of HCI are largely understood, but the key variables for effective H/E teamwork are not yet described as analytical criteria. Some parameters of trust have been identified (Muir, 20).

3) Development of instruments and piloting techniques for measurements in flight, and a deliberate practical scheme for research. Some work has been done on measuring Teamwork (Sherwood Jones et al (17), the work of Taylor, and the High Reliability Organization programme (6, 7)), but there are no programmes to the author's knowledge that are expressly addressing the instrumentation needs of H/E team design and assessment.

4) Growth and refinement of pilot/operator opinion regarding desirable qualities. The limited exposure that operators have had has been fragmented, and no pool of user opinion has been allowed to build up.

5) Measurement of characteristics and assessment of results for a cross section of applications. This has still to get under way.

## CONCLUSIONS

Vincenti (4) divided the process of maturing understanding into two stages; the establishment of

basic analytical understanding and practical capability, and putting this understanding and practical capability to work to establish concrete workable specifications. From the above sketch, it would appear that the H/E Crew is still in the first stage of maturity and in need of collective empirical experience.

There has been considerable change in the context (since 8, 21) in which standards for H/E teamwork need to be set. The most promising way forward at this time is to develop a checklist approach to designing and assessing teamwork under a safety management umbrella. This paper has outlines a framework that could be developed to form a checklist for design guidance and for test and evaluation of knowledge systems with human-electronic components.

## ACKNOWLEDGEMENTS

## REFERENCES

1 Hollnagel, E., Woods, D.D., 'Cognitive Systems Engineering: New Wine in New Bottles',Int.J.Man-Machine Studies, 18, pp 583-600, 1983

2 Woods D.D., 'Cognitive Technologies: the Design of Joint Human-Machine Cognitive Systems' AI Magazine 6, no 4, 1986

3 Newell, A., 'Unified Theories of Cognition', Harvard UP, 1990

4 Vincenti, W.G., 'What Engineers Know and How They Know It', John Hopkins UP, 1990

5 ISO 9241, Ergonomic Requirements for Office Work with Visual Display Terminals (VDTs). Part 2. Guidance on Task Requirements, part 10. Dialogue Principles, Part 11. Guidance on Usability Specification and Measures

6 Rochlin, G.I., 'Technology, Hierarchy and Organizational Self-Design US Naval Flight Operations as a Case Study', Working Paper 88-18, Institutes of Governmental Studies, Univ Cal Berkeley

7 LaPorte, T.R., Consolini, P.M., 'Working in Practice by Not in Theory: Theoretical Challenges of "High Reliability Organizations"', J.Public Admin. Research and Theory, 1, pp 19-47, 1991

8 Reising, J., Emmerson, T., 'Standardisation Issues for the Year 2000', Contemporary Ergonomics 1986, Proceedings of the Ergonomics Society Annual Conference, Taylor and Francis, 1986

9 Sherwood Jones, Salway, A., Ferguson, N, Holt, J., 'A method for the Assessment of Command System Organizations' DRA Maritime, TDP/34/REP/0001, June 1992 (UK RESTRICTED)

10 Sherwood Jones, B.M., Holt, J., Ferguson, N., Narborough-Hall, C.S., Brennen, S.D., 'A Theory of C$^2$ Organizations" IDASCO Conference, IEE CP 353, 1992

11 Def Stan 00-56, 'Safety Management Requirements for Defence Systems Containing Programmable Electronics'

12 Carr, D.J., Salway, A.F.S., Sherwood Jones, B.M., 'Safer by Design: A Framework for a Design Process Standard for vehicle MMI', ISATA 1993, Florence

13 Quinn, R.E., Rohrbaugh, J., 'a Spatial Model of Effectiveness Criteria: Towards A Competing Values Approach to Organizational Analysis', Management Science, 29, no 3, 1983

14 Rohrbaugh, J., 'Assessing the Effectiveness of Expert Teams', NATO ASI Series Vol F35, Expert Judgment and Expert Systems ed. Mumpower, J. Springer-Verlag 1987

15 Silverman, B.G., 'Critiquing Human Error', Academic Press, 1992

16 Morgan, M.G., Henrion, M., 'Uncertainty' CUP, 1990

17 Sherwood Jones, B.M., Northcote, E.A., 'TDS User Knowledgebase Interface Validation' DRA Maritime, TDP/6.11, 1991 (UK RESTRICTED)

18 Edelson, T., 'Can a System be Intelligent If It Never Gives a Damn?', Proceedings AAAI Fifth National Conference on AI, Philadelphia, 1986

19 Sloane, S.B., 'The Use of Artificial Intelligence by the United States Navy: Case Study of a Failure' AI Magazine, 12, no 1, 1991

20 Muir, B. E., 'Trust between Humans and machines, and the Design of Decision Aids' Int. J.man-Machine Studies, 27, pp 527-539, 1987

21 Sherwood Jones, B.M., 'Standardisation Needs for Interaction with Intelligent Systems', IEE Colloquium on 'User Interfaces and Standardisation' 27 Jan 1989

# SESSION V - SYSTEMS INTEGRATION LESSONS

## SYNOPSIS

The final section reports lessons learnt from evaluations of HE-C concepts in relatively mature programmes. The papers provide updates on National projects reported at previous HE-C meetings. Paper 23 reports the status of the German Cockpit Assistant System, CASSY; Paper 24 reports the evaluation of the British Mission Management Aid (MMA); Papers 25 and 26 develop ideas from the USAF Pilot's Associate (PA) programmes; Paper 27 reports the latest phase of the French "Copilote Electronique" project. *Paper 23* provides a description of the role and functioning of CASSY, a knowledge-based pilot support or assistant system. CASSY has recently been integrated into the cockpit of a test aircraft, where it has undergone 11 hours of flight testing. CASSY comprises a dialogue manager for controlling the interface information flow, including speech input/output. Other KBS modules include automatic flight planning, piloting, pilot intent error recognition, monitoring flight status, environment and systems, and execution aiding. Test flights with two pilots, and additional airline pilot observers, included operating through high density air traffic at major German airports. The flight tests successfully proved CASSY's functioning. Speech control was used for complex inputs. The autonomous flight plan functions compared favourably with current flight management systems. CASSY warnings and prompts were generally considered to be justified, and frequently initiated appropriate corrective action. Aiding situation assessment is judged to be the key benefit of KBS. *Paper 24* describes the development of the man-machine interface of the UK MMA demonstrator, which uses KBS for some limited functions. Key activities for the MMA were identified as management of position, time, fuel, and EW. Functional assessments were undertaken by nine aircrew in a simulation of a ground attack mission. The paper reports the level of crew acceptance of the MMA tasks, levels of authority and display formats. Concern was expressed about MMA performance of some critical tasks (defensive options, automatic weapon aiming), and about MMA responsibility for position and EW management. High integrity sensor information would be needed for the MMA decision processes to be considered valid. *Paper 25* characterises the "associate relationship" as a mixed initiative approach to collaborative problem solving. The authors describe how lessons learnt from the PA program have led to the development of machine learning tools, re-using associate knowledge bases, and associate applications in novel domains, including supporting multiple operators. In the Rotary Pilot's Associate programme, it has been found that aiding helicopter operations involves decreased sequential behaviour, compared with fixed wing missions, and hence more difficulty in providing context sensitive plans. Research on how and when humans will decide to use an associate indicates that workload does not affect automation reliance, whereas task complexity, automation reliability and fatigue are important factors. Trust and automation state uncertainty contribute to automation use decisions. The wide range of individual differences in automation use strategies leads to different under- and over-use tendencies. Commercial pilots are strongly biased to automation use, even through failure periods. The effects of inappropriate reliance on system safety need to be better understood before systems are fielded. *Paper 26* describes the development of principles of interaction, equivalent to prime directives, under which intelligent systems are given authority to take actions. The principles need to be based on measurements that can be estimated concurrently with, or predicted in advance of, the exercise of the authority, such as workload. Global principles for EC functioning include: ultimate human control and over-ride authority; always following the human's lead; monitoring and supporting the human and not the opposite; never seizing control; behaviour always being predictable; always being able to be turned off. Specific authorisations are derived from global principles by iteratively incorporating contextual details within the rules. A hazard monitoring system is discussed as a specific application of this rule derivation process. The authors consider that provision of a human factors of systems functionality would be a major step towards predictable, consistent and trustworthy HE-C systems. *Paper 27* relates the status of the Copilote Electronique project, at the start of an exploratory development aimed at non real-time simulation of Rafale SU2 missions. Design difficulties are reviewed, and a functional analysis of KBS is presented as a model for pilot assistance. Situation assessment, plan management, co-ordination, and plan generation are considered to be the main KBS domains. The authors summarise a set of user-oriented ergonomics rules intended to guard against KBS applications failure. A rapid prototyping life-time model is proposed, with iterating, evolutionary characteristics for integration of ergonomics compliance. Pilot meta-knowledge is identified as the key to mission planning and reflex actions. Conventional KA techniques are inadequate for eliciting this knowledge, and new methods have had to be developed. Providing an exchange language for common plans and goals, and intent recognition planning, are particularly difficult technical areas.

# CASSY - The Electronic Part of a Human-Electronic Crew

**M. Gerlach, R. Onken**
Universität der Bundeswehr München
D-85577 Neubiberg
Germany

## 1. SUMMARY

This paper describes the knowledge-based Cockpit Assistant System CASSY and its functions as an example of human-centered automation. The paper reminds of the requirements for human-centered automation in terms of the two basic requirements which the numerous known requirements can be boiled down to easing the task of technical mechanization. These requirements are the basis for the CASSY-design. Finally, the results of CASSY are given.

## 2. PREFACE

Despite the introduction of many new technologies in modern aircraft cockpits, pilot error is still the major cause of aircraft accidents today. Investigations show that approximately 75% of aircraft accidents are due to human error. In most cases the pilots were overburdened with a situation which they could not cope with because of their natural human limitations and weaknesses. Typical limitations are the constrained capacity of human short term memory and a confined information gathering process. Even human capabilities which usually exhibit extraordinary performance may also be the reason for human failure. The usual experience of success in pattern recognition may mislead to uncritical behaviour and result in human error since crucial cues may be ignored. In aircraft accidents these limitations resulted in faulty situation awareness and interpretation also leading to unsufficient planning and decision-making. The problem of a lack of situational awareness and its possible fatal consequences is pointed out in [Endsley, 92].

Automation was considered to be a promising approach to solve the problem. In order to reduce workload and to overcome human weaknesses a number of tasks have been taken over by the machine. Those functions were assigned to the machine which could technically be automated, the others remained with the pilots. Therefore, current automatic flight guidance systems like autopilots or flight management systems support on the skill-based or only partly on the rule-based level of cogni-

## Flight Guidance and Control Today



Figure 1

tive human operator behaviour. Automation on the high level of knowledge-based behaviour appeared to be too difficult to be realized. Unfortunately, technical aspects have been the major driving factors of automation instead of the pilots' needs [Hollnagel, 93].

Today's task distribution in aircraft cockpits, in particular transport aircraft, is depicted in figure 1. Electronic support covers the fields of sensorics, flight plan calculations, and functions for flight plan execution. As the systems do not have the knowledge to understand the underlying situation, extensive communication effort is needed to inform the system of the situation knowledge (thick arrow in figure 1). Modern flight management systems, for instance, require both data collection and situation assessment by the pilot and a difficult man-machine interface to ensure that the system uses the appropriate data.

As a result of the level of cockpit automation, the interaction between the crew and the machine has become more and more complex. The pilots are not always sure of the current state of their systems [Wiener, 89; Dycke, 93] . This led to new forms of pilot error [Chambers, 85].

## 3. BASIC REQUIREMENTS FOR AUTOMATION

Human-centered automation [Billings, 91] is a new approach to achieve a better design of man-machine systems. Design criteria were established which aim at a cooperative function distribution between man and machine like that of two partners [Reising, 93]. Both man and machine are active at the same time, in contrast with today, and assist and control each other while heading for the same goals.

As far as flight safety is concerned, successful automation should be designed to avoid situations in which the crew is overburdened, i. e. crew demands and workload should be kept on a normal level for any situation and its corresponding tasks like situation assessment, planning, and plan execution. Derived from this demand on automation an electronic pilot support system should comply with two basic requirements [Onken,93; Onken, 94]:

- 1.) Within the presentation of the entire flight situation the system must ensure to guide the attention of the cockpit crew towards the objectively most urgent task or subtask.

- 2.) If the aforementioned requirement is met, and if there (still) occurs a situation in which the cockpit crew is overburdened then the situation has to be transformed - by use of technical means - into a situation which can be handled normally by the cockpit crew.

The first of the two requirements is concerned with the aspect of keeping up the pilots' situational awareness. But it also implies that the support system itself performs a correct situation evaluation which is necessary to provide intelligent automated functions. If the system is not able to understand the underlying situation its functions might work on the basis of wrong assumptions. To date, this is the reason why functions of flight deck automation operate in unsufficiently intelligent and sometimes dangerous manner. As there was no technology available for situation assessment by machine, automation was limited to the second half of the second requirement and the other aspects were ignored.

Improvements can be expected from new technologies and methods in knowledge processing and artificial intelligence. Examples of knowledge-based computer systems show that these technologies mature.

## Flight Guidance and Control Upcoming



Figure 2

A real partnership between pilot and electronic copilot in the sense of the two basic requirements is established when the support system is deployed in a situation-dependent and versatile task distribution as proposed in figure 2. Compared with figure 1 the interactions and activities of the two partners take place on more and different levels of cooperation. High level functions are added to the machine which enable it to understand flight situations and to consider goals, subgoals and plans on its own.The functional capabilites are implemented in the system in parallel to the cockpit crew. Like a cooperative crew the two influence each other but are also able to work independently. As long as the crew is fully responsible for the flight the decisions and in-flight selection of goals are assigned to the pilots. At the time being, the knowledge based assistant system is far from holding any final decision authority. Its merit is that it takes all amenable information into account to interpret a given situation whereas pilots decide on fragmental excerpts of situational information due to limited sensing and processing capacity. Computer support systems analyze a situation on fixed strategies which they will repeat exactly if the situation recurs exactly and which might fail in unforeseen situations. Especially those situations could possibly be mastered by the human pilot because of his broader learning horizon.

The interaction on the different levels of cooperation puts new demands on man-machine dialogue. The complexity and versatility of data to be trans-

mitted increases with the integration of high level functions. The first basic requirement calls for a communication design in which crew resources are considered to select an appropriate means, and to focus the crew's attention on the right task. It must be sure that the crew percieves all relevant information. This can only be achieved if the information flow to the pilot is managed in a centralized way taking into account all knowledge of the current situation. The use of natural communication forms like speech seems to be a good approach to improve human-computer dialogue.

## 4. THE COCKPIT ASSISTANT SYSTEM *CASSY*

With the following description of the Cockpit Assistant System CASSY, we would like to present an example how to design to comply with the discussed ideas. CASSY was developed and flight tested at the *Universität der Bundeswehr München* (UniBwM) in cooperation with *DASA-Dornier*.

In the previous chapter the important part of electronic situation understanding for successful machine support was pointed out. A system can only understand a situation if it has the appropriate knowledge of the problem space it works in. Since CASSY is limited to civil aviation, its knowledge base comprises the elements of figure 3. This knowledge base is characterized by static knowledge, e.g. a normative model of cockpit crew behaviour or knowledge of the used aircraft, and dynamic knowledge referring to in-flight changing

Figure 3

circumstances caused by instructions from air traffic control (ATC) or environmental influences. Stored in a central situation representation, this knowledge serves as a picture of the current situation.

Figure 4

In order to gather dynamic knowledge and to transmit its conclusions the assistant system is placed in the flight deck as indicated in figure 4. CASSY has interfaces to the flight crew, to the aircraft, and to ATC. The interfaces ensure that all knowledge sources are available for the task specific modules of the system which are shown on the left column of figure 5.

Figure 5

The **Automatic Flight Planning** module (AFP) generates a complete global flight plan [Prevot, 93]. On the basis of its knowledge of mission goal, ATC instructions, aircraft systems status, and environmental data an optimized 3D/4D trajectory flight plan is calculated. The flight plan, or several plans, is presented as a proposal which the crew accepts or modifies. Once a flight plan is chosen it serves as a knowledge source for other CASSY modules. The AFP recognizes conflicts which may occur during the flight, e.g. due to changing environmental conditions or system failure, and appropriate replanning is initiated. If necessary, this re-

planning process includes the evaluation and selection of alternate airports. Since the module has access to ATC instructions, radar vectors are incorporated in the flight plan autonomously and the system estimates the probable flight path ahead.

The presentation of the resulting situation-dependent flight plan to the crew serves directly the realization of the first requirement discussed in chapter 3 because the consequences of the valid flight plan is obvious to the crew. The extensive aid in decision-making and time consuming flight plan calculations supports the second requirement, too.

The Module **Piloting Expert** (PE) uses the valid flight plan to generate necessary crew actions. It is responsible for processing a crew model on normative and individual crew behaviour [Ruckdeschel, 94]. The normative model describes the deterministic pilot behaviour as it is published in pilot handbooks and air traffic regulations. The model refers to flight guidance procedures concerning altitude, speed, course and heading, but also to aircraft systems management. Given the flight plan and a pointer on the current leg, provided by the Monitor of Flight Status, the system determines the appropriate normative values and tolerances on aircraft systems and flight status data. Using the individual model, determined from an adaptive component, these data are adjusted to individual preferences.

The crew model as used to generate necessary crew actions, is absolutely vital to meet requirement 1. It enables the system to identify the most important actions on the basis of the underlying situation and to interpret the observed crew behaviour.

The expected crew actions are compared with the shown behaviour of the crew in the module **Pilot Intent and Error Recognition** (PIER) [Wittig, 93]. The crew actions are derived indirectly by interpreting the aircraft data. If given tolerances are violated, the crew will be informed by hints and warnings and the detected mistake is pointed out to the pilots. In the case the crew deviates intentionally from the flight plan, the module checks if the shown behaviour fits to a given set of intent hypotheses which are also part of the crew model. These hypotheses represent behaviour patterns of pilots in certain cases, e.g. tasks to be done when commencing a missed approach procedure or to deviate from the flight plan to avoid a thunderstorm ahead. When an intentional flight plan deviation and the respective hypothesis is recognized, appropriate support, e. g. replanning, is initiated.

The monitoring of the pilots' actions and the distinction between error and intentional behaviour in extraordinary situations serves both basic requirement 1 and 2.

Additional monitoring modules are needed to enable the system to recognize and interprete current situations. The **Monitor of Flight Status** provides the present flight state and progress. It is also able to report the achievements of subgoals of the flight. The **Monitor of Environment** gathers information of the surrounding traffic, e.g. from TCAS and of weather conditions, and incorporates a detailed navigational data base of the surrounding area. The health status of aircraft systems are monitored by the **Monitor of Systems** like a diagnosis system.

Obviously, the monitoring systems are essential to meet the first requirement as their outputs are an important part of the full picture of the present situation. Since their output is also used to adjust the flight plan to the situation, they contribute to meet the second requirement, too. Additionally, the continous observation of flight progress, environment, and aircraft systems supports the crew on tedious tasks.

Communication plays an important role in CASSY. The kind of information to be transmitted in either direction varies with respect to the different modules (figure 6). The information flow from CASSY to the crew and vice versa is controlled by the module Dialogue Manager [Gerlach, 93]. The many different kinds of messages require a processing in order to use an appropriate display device and to present the message at the right time. As output devices both a graphic/alphanumeric colour display and a speech synthesizer are used. Short warnings and hints are used to make the crew aware of a necessary and expected action and are



Figure 6

transmitted verbally using the speech synthesizer. An additional alphanumeric line is fixed on the graphic display to facilitate perception of difficult verbal messages. More complex information, e. g. the valid flight plan, is depicted on a moving map on the graphic display.

Another important feature of the DM is that since the tolerances and danger boundaries are given in the crew model and the necessary actions are inferred, a priority ranking of the output message is evaluated and the most important message is issued with priority.

The input information flow is established by use of speech recognition in addition to conventional input mechanisms. In order to improve speech recognition performance, almost the complete knowledge of CASSY is used to provide situation depending syntaxes. Thus, the complexity of the overall language model is reduced significantly. Not only the pilot's inputs must be considered but also the inputs from ATC. The data link, indicated in figure 6, is not available to date. Discrimination of ATC instruction from pilot input is achieved by picking up the pilot's verbal acknowledgement of the ATC controller's instructions.

The use of speech input and output devices also reflects the idea of a human-electronic crew and of cooperation of partners.

In figures 6 and 7 an additional module is shown which is called **Execution Aid** (EA). In this module several functions are realized which can be called by the crew. Aircraft settings, navigational calcula-



Figure 7

tions and data base inquiries are carried out. These functions are similar to available automated functions in today's aircrafts and are mainly designed to meet requirement 2. For the pilots, the main difference is the use of speech input which facilitates the use of these services.

The integrated system, its modules, and its control flow is given in figure 7.

## 5. RESULTS OF THE FLIGHT TESTS

In June 94, CASSY has undergone an eleven hours flight test campaign in Braunschweig, Germany.

The modules of CASSY have been implemented on an off-the-shelf available *Silcon Graphics Indigo* workstation using the programming language *C*. A *Marconi MR8* PC card was used as speaker dependent, continous speech recognition system. A *DECTalk* speech synthesizer served as speech output device using three different voices enabling the pilot to discriminate between the various messages. The components were connected using serial lines and ethernet.

The system was integrated into the test aircraft *ATTAS* (Advanced Technologies and Testing Aircraft) of the *Deutsche Forschungsanstalt für Luft- und Raumfahrt* (DLR) in Braunschweig. The aircraft is well equipped for flight guidance experiments as it is possible to operate the aircraft via a single seat, experimental cockpit located in the cabin. An ethernet connection to the CASSY workstation was used to simulate an avionic bus system as aircraft interface in either direction. As ATC interface both approaches were tested a simulated ATC data link and the pilot's acknowledgement of ATC instructions.

The test flights comprised instrument flights from the regional airport Braunschweig to the international airports of Frankfurt, Hamburg and Hannover at which a missed approach procedure was conducted before returning back to Braunschweig.

The experiments proved CASSY's functions from take-off to landing throughout the complete flight. Speech recognition performed well in the aircraft as the surrounding noise was primarily engine noise which did not change much during flight. The recognition rates were similar to those achieved in the more quiet simulator environment at the University in Munich where CASSY was developped and tested before.

One important aspect of the tests was to prove the system in the high density air traffic control of German airports which could not be tested in the

scope of simulator test runs. During the campaign, any given ATC instruction could be processed and integrated into the flight plan by CASSY. Compared to available flight management systems the autonomous integration of ATC radar vectors proved to be faster and did not lead to distracting information input.

On the basis of the flight plan the correct expected pilot actions were generated and pilot errors, provoked or non-provoked, were detected and the appropriate warnings were issued. Wrong warnings occured infrequently and were uncritical in any case.

Two pilots were flying with CASSY in the test aircraft. Additional pilots from Lufthansa German Airlines were participating to observe the tests and to take part as a second pilot aside the test pilot.

CASSY was well accepted by the pilots throughout the campaign. In particular, the pilots appreciated the autonomous flight plan functions of CASSY. Warnings and hints were considered justified and corrective system inputs were made. Speech input was generally used when complex inputs were to be made, e. g. frequency settings which could be made using the simple name of the station instead of its difficult frequency.

## 6. CONCLUSION

Today's operational flight deck automation has come to a state in which pilot-aircraft interaction becomes more and more difficult. The reasons why recent automation did not lead to a reduction of pilot induced accidents can be found in the insufficient capability of the automated functions to recognize and interpret the underlying situation in a full picture and thereby complying with the basic requirements for human-centered automation. Therefore, adequate situation assessment is a key function for successful machine support. Knowledge-based systems making use of methods of artificial intelligence and knowledge processing seem to be a promising approach to provide this key function.

The Cockpit Assistant System CASSY is an example of how a pilot support system could look like. It is designed to meet the basic requirements for machine support. The successful flight test of the system shows that a new generation of this kind of electronic automation can be introduced in the mid-term future.

## 7. REFERENCES

Endsley, 92
Endsley, M.R.; Bolstad, C.A.
"Human Capabilities and Limitations in Situation Awareness"
AGARD Joint FMP/GCPSymposium, Edinburgh, 10/92

Hollnagel, 93
Hollnagel, E.; Warren, C.; Cacciabue, P.C.
"Automation in Aerospace and Aviation Applications: Where are the Limts?"
Workshop Report from 4th HMI-AI-AS, Toulouse, 09/93

Wiener, 89
Wiener, E.L.
"Reflections on Human Error: Matters of Life and Death"
Proceedings of the Human Factors Society, 33rd Annual Meeting, 1989

Dyck, 93
Dyck, J.L.; Abbott, D.W.; Wise, J.A.
"HCI in Multi-Crew Aircraft
Proc. of theHCI, Orlando, 19 A, 08/93

Chambers, 85
Chambers, A.B.; Nagel, D.C.
"Pilots of the Future: Human or Computer?"
Communications of the ACM, Vol. 28, No. 11, 1985

Billings, 91
Billings, C.
"Human-Centered Aircraft Automation: A Concept and Guidelines"
NASA Tech Memorandum 103885, Moffet Field, 1991

Onken, 93
Onken, R.
"Funktionsverteilung Pilot-Maschine: Umsetzung von Grundforderungen im Cockpitassistenzsystem CASSY"
DGLR-Tagung d. Fachauss. Anthropotechnik, Berlin, 1993

Onken, 94
Onken R.; Prevot, T.
"CASSY - Cockpit Assistant System for IFR Operation"
19th ICAS - AIAA Aircraft Systems Conference, Anaheim, CA, USA, 18 - 23 Sept. 1994, in press

Prevot, 93
Prevot, T.; Onken, R.
"On-Board Interactive Flight Planning and Decision Making With the Cockpit Assistant System CASSY"
4th HMI-AI-AS, Toulouse, 1993

Ruckdeschel, 94
Ruckdeschel, W.; Onken, R.
"Modelling of Pilot Behaviour Using Petri Nets"
15th Int. Conf. on Application and Theory of Petri Nets, Zaragoza, 1994

Wittig, 92
Wittig, T.; Onken, R.
"Pilot Intent and Error Recognition as Part of a Knowledge Based Cockpit Assistant"
AGARD Joint FMP/GCP Symposium, Edinburgh, 10/92

Gerlach, 93
Gerlach, M.; Onken R.
"A Dialogue Manager as Interface Between Aircraft Pilots and a Pilot Assistant System"
Proc. of theHCI, Orlando, 19 A, 08/93

Reising, 93
Reising, J.M.; Emerson, T.J.; Munns, R.C.
"Automation in Military Aircraft"
Proc. of theHCI, Orlando, 19 A, 08/93

# MODELLING THE INFORMATION FLOW - DEVELOPMENT OF A MISSION MANAGEMENT AID FOR FUTURE OFFENSIVE AIRCRAFT

By

**J.M Davies**
**British Aerospace Defence Ltd**
**Military Aircraft Division**
**Farnborough, Hants, UK**

## SUMMARY

*The trend towards more complex avionics systems integrated into high performance aircraft, coupled with an uncertain threat environment, presents a potentially intolerable workload situation for the Pilot in the next generation single seat cockpit.*

*To address the information 'battlefield' a 5 year collaborative programme between the 4 major UK Military industrial organisations (British Aerospace, Ferranti, GEC Avionics, GEC Sensors and Smiths Industries) and the MOD RAE (now Defence Research Agency) was set up in 1988.*

*The Primary objectives were to investigate the concept of an intelligent Mission Management Aid (MMA) and to develop a set of functional requirements. The four main functional areas selected for examination were Sensor Fusion, Situation Assessment, Dynamic Planning and the Man-Machine Interface.*

*This paper describes the development of the Man-Machine Interface with particular emphasis on the cognitive interface. The approach taken was to consider the teaming aspect of the man and the machine in the execution of various on-board management functions.*

*The initial on-board management functions chosen were Position Management, Fuel Management, Time Management and EW Management. Each management function was decomposed into a series of tasks and task modelling was undertaken within the MMA in order to establish the decision processes and allocation of responsibilities associated with the execution of these tasks.*

*A Proof of Concept Simulation (PCS) rig was developed in order to evaluate the MMA functionality and nine aircrew with a wide range of platform experience participated in a final assessment immediately prior to the closure of the project. There was a fairly consistent response on the operational utility of the integrated MMA and their comments are included in this paper.*

# 1. MMI DESIGN APPROACH

In order to achieve a high level of Human Machine Integration in the development of the Mission Management Aid (MMA) a different approach to the standard allocation of function phase was adopted.

The approach taken was to consider the teaming aspect of the man and machine in the execution of the various on-board management functions.

In order to achieve this there were several key steps which are described in this paper:
- Identification of key management activities
- Development of model mission
- Development of detailed task models
- Integration into Proof of Concept Rig
- Aircrew Assessments

# 2. IDENTIFICATION OF KEY MANAGEMENT ACTIVITIES

The key management activities were identified through a fairly intensive knowledge elicitation exercise which involved discussions with SAOEU Boscombe Down, Flight Operations BAe Warton, RAF Marham and RAF Coltishall. It was evident from the descriptions of the tasks carried out that there was an executive set of management functions associated with prosecuting the mission. The management functions identified initially were:
- Position Management
- Time Management
- Fuel Management
- EW Management

Whilst there is a strong interdependency between these management functions the aim of the exercise was to identify the tasks relating to the functions in isolation and examine the decision process and likely allocation of responsibilities between the team members (man and the machine).

## POSITION MANAGEMENT

Position Management broadly entails knowing where you are, where you want to get to and how to get there. In knowing how to get there a knowledge of the external world in terms of

terrain features (natural and man-made), threat situation and NO-GO zones is required.

The MMA needed to address the following issues:
- When is a re-route situation imminent
- How to assess possible re-route solutions
- How to communicate with the Pilot in regard to re-route situations

## TIME MANAGEMENT

Time Management broadly entails ensuring accurate Time On Target (TOT) and waypoints and currently includes time related tasks such as manual IFF updates. The key issues for this function were as follows:
- To identify the type and level of information relating to time management tasks
- To identify the situations where time related tasks are carried out (e.g. dealing with revised time on target, updating IFF codes)
- Ensuring accurate Time on Target

## FUEL MANAGEMENT

Fuel Management broadly entails ensuring that there is adequate fuel, that it is sequencing correctly at the appropriate rate and taking the corrective action where necessary. The issues to be dealt with for this management function were:
- To determine whether there is sufficient fuel to complete the mission and land with the predetermined fuel load
- Determine how to deal with a low fuel situation (optimum flight profile, in-flight re-fuelling
- Dealing with information requirements - bingo fuel (predicted fuel usage) combat fuel (continguency fuel) chicken fuel ( Get-you-home fuel)

## EW MANAGEMENT

Electronic Warfare Management broadly entails monitoring the threat situation (air and ground) and taking the necessary counter actions. The key issues for EW management were:

- To identify the level and prioritisation of information required to interpret the threat situation
- To identify the dialogue requirements associated with interrogating the threat information
- To determine what evasive actions to take

## 3. MODEL MISSION

A ground attack model mission was developed to exercise the embodied functionality. It contained a series of canned events covering internal system failures, re-tasking, and unexpected hostile activity.

The mission profile was centred about the central European arena with Munster as the home base. A Forward Line of Own Troops (FLOT), Fire Support Co-ordination Line (FSCL) and IFF line were used in conjunction with Air Co-Ordination Order (ACO) corridors in friendly territory. The target was a power station in the Hartz mountains and an initial route was planned using the best terrain and designed to avoid known SAM sites . This baseline route with waypoints, speeds and intelligence data was provided to the MMA. A description of the scenario is given in Section 5.

## 4. TASK MODELLING

Prior to defining the task models, each of the management tasks were analysed in terms of the questions relating to the management activities for which answers are required. In this way the machine should be able to anticipate the need for information or action.

A task may be described in simple terms by:

- The stimulus to perform the task
- The information required
- The decision process within the task
- The task output

Thus the tasks were identified in terms of their goals, stimulus, decision-base, actions and output. All tasks were deterministic, rule based and event driven.

An example of a typical task is ACHIEVEMENT_OF_TIME_AT WAYPOINT.

The task goal is to determine whether the time at next waypoint is achievable within the speed and fuel constraints.

The information required to perform task - DEMANDED_SPEED, SPEED_CONTSTRAINT , PREDICTED_FUEL_RATE will be output from other tasks.

The decision process will determine whether the speed constraints is satisfied - if DEMANDED_SPEED > SPEED CONSTRAINT then a warning to the pilot will be activated together with the ADVISE_NEXT_WAYPOINT task.

If the speed constraint is not violated then PREDICTED_FUEL task will be activated substituting DEMANDED_SPEED for the current planed speed value.

If COMBAT_FUEL <0 then a warning to Pilot is activated and ADVISE_ NEXT_ WAYPOINT task is activated.

This decision process mirrors the mentally intensive process currently used by Aircrew. This particular example covers Time, Position and Fuel Management tasks.

## 5. SCENARIO DESCRIPTION

The following scenario was used during the assessments:

The aircraft took off from Munster airfield and transited out of Base Defence Zone.

Whilst transiting between waypoint 3 and 4 two system problems were identified and communicated to the Pilot; the first related to a Fuel transfer failure and the second related to a degradation in the Terrain Referenced Navigation system. In the case of the second problem the MMA suggested a change of Set Clearance Height (SCH) to 1000ft which was accepted.

Whilst still within the friendly ingress phase a 3 minute delay to TOT was received and the MMA

offered an orbit solution to lose the time . There are several other options including abort mission but the orbit solution was chosen.

Shortly after completing the orbit the TRN system was fully functional again and the MMA recommended descending to the planned height.

As the Aircraft crossed the defined IFF line into hostile territory the IFF was automatically switched to standby.

A dynamic threat assessment module advised on the feasibility of the planned route based upon a prediction of the SAM launch opportunities. It was possible to evaluate the effect of aircraft height changes upon the SAM engagement zones and thus assist in the position management task. During the hostile ingress phase there were several instances of re-route due to acquisition by SAMs. As the aircraft approached the IP the weapons were automatically armed.

Within the attack phase the aircraft was acquired by a SA13 but re-routing was not possible due to time constraints. Defensive options were automatically employed with the dispensing of chaff and flares.

As the aircraft egressed through the hostile territory there were several occasions where the aircraft is acquired by ZSU 23/4 and communicated to the Pilot. The necessary evasive action was taken.

As the aircraft transited the FLOT a low fuel warning is declared and alternate airfield located within fuel range. An alternative option is to carry out in-flight re-fuelling
and route to tanker was also available.

## 6. PROOF OF CONCEPT RIG DESCRIPTION

The rig was configured as shown in figure 1. The majority of the MMA functionality was contained within the Symbolics and Sun work stations whilst the displays management was controlled within a Silicon Graphics environment.

There were two display surfaces containing a Pilot Awareness display and Aircraft Status displays

**Pilot Awareness Display**

The Pilot Awareness display (see figure 2) comprised various optional views of the external scene including a two dimensional digital map , perspective view and a Gods Eye view.

It was possible to add, move or delete waypoints as required and the MMA could keep track of the changing scenario advising on feasibility of the route in terms of time, fuel and threat constraints. The multiple threat picture based upon fused information from the sensor fusion module was overlaid on the map and the least risk route calculated by the tactical planner module displayed as a route vector line approximately 10 kilometres in length . The predicted missile engagement zones were depicted in red on the route line for areas of high threat

A cross-sectional view of the route was provided to allow assessment of the vertical clearance the aircraft had when flying under a threat zone.

This display also contained an MMA status display containing the various options that the MMA had to communicate to the Pilot. Different classes of messages were included within the scenario:

- Warning of high threat and advise re-route (re-route options on main display with step function to allow access to different re-route solutions
- Information messages (IFF set to Standby/on, Bingo Fuel States)
- Warning of system problem plus advised action (TRN degraded - SCH 1000 ft)
- Advisory message of in-cockpit task completed ( EW Function check complete)

The routing options could be accepted or rejected as appropriate. It should be emphasised that this type of display would not be appropriate for an airborne application but merely served as a communication channel for demonstration purposes. DVI and DVO would be a more suitable implementation for this type of cockpit function.

## Status Display

The second display contained various aircraft status information (Fuel, Weapon, Radar Warning Display and Missile Approach Warner) and a standard Fast Jet HUD Display (see figure 3).

The Fuel displays allowed the information from the fuel management module to be displayed and included predicted fuel to complete mission (based upon the pre-planned route), combat fuel, bingo fuel states, dynamic chicken fuel state and fuel to tanker.(see figure 4). If combat fuel reached zero then tanking or diversion options were presented on the MMA dialogue display. The status display also contained a risk display which was intended to provide the machine's interpretation of the threat situation.

## 7. ASSESSMENT RESULTS

Nine aircrew participated in the assessment which was more accurately a demonstration of the MMA functionality within a work-station environment. They were briefed on the objectives of the session which were as follows:

- To examine the operational utility of the MMA functionality
- To determine whether they were in agreement with the solutions that the MMA had identified and whether the MMA decision process mapped onto their own
- To identify the level of dialogue with the MMA that would be considered necessary and the consequential level of authority that the MMA should be given
- To investigate the information requirement - is it in the right form at the right time to assist the decision process

After the assessment/demonstration the crews completed individual questionnaires followed by a discussion of the broad conceptual issues. The concept of task allocation based upon existing aircrew activities was explained and they were asked to comment upon this allocation against three criteria:
- Is it a suitable machine task - Do you feel comfortable with the MMA performing that task

- Was the task performed satisfactorily
- What level of authority should be given to the machine
  1) = Do it and do not inform
  2) = Do it and provide feedback
  3) = Don't do it until instructed.

## Summary of Results

1. In general the majority of the aircrew considered that the tasks were suitable for the machine to perform . There was concern by three of the aircrew on the defensive options tasks and by two on the automatic weapon arming.

2. Most of the tasks were completed satisfactorily although there was some concern regarding the decision processes associated with the defence options. Two of the aircrew expressed concern about the re-routing in that it did not take account or have access to all the factors that they would normally use (e.g. weather).

3. The level of authority for all the basic monitoring tasks were seen as level 1. Eight crew concurred that auto re-routing was a level 3 task and one of the crew thought that both the attack option and defensive options should be on selection only (level 3). The rest of the tasks were identified as level 2 .

4. In terms of the information content on the displays the crews had the following observations to make:

 - The intervisibility display had high operational utility especially when used in conjunction with the 2 D slice display

 - The predicted missile engagement zone on the route line also had high utility value

- The perspective displays with the 3D threat envelopes were considered to have limited value

- The Gods Eye View display was considered to be a useful display for ground based mission planning use but was considered unsuitable for an airborne environment

- The risk display was considered to be open to interpretation and ambiguous and therefore was seen as having low utility value.

- When questioned about the categorisation of the threat into type specific (SA8,SA11 etc.), short, medium or long range, or Radar/IR guided all aircrew thought that radar/IR guided was the most useful categorisation

- All crews felt that there was adequate information displayed to assess the threat situation and choose a re-route option. However they were of the opinion that it would be preferable to manually change the route and for the MMA to dynamically assess the route using time, fuel and threat costs.

## 8. CONCLUSIONS

Many of the comments related to the presentation of the information and it will be the ability of the crew to rapidly interpret and react to the intent of the MMA that will be the key to the successful integration of this type of system.

The attempt to create a communication channel display for the MMA on the demonstration, whilst being contrived, did serve to demonstrate the danger of overloading the man with interrupts. Man is essentially a single channel processor and it was evident from the assessment that further work was required to prioritise the interrupts and protect him from the system, particularly at crucial phases of the mission.

There was immediate crew acceptance for the machine to take responsibility for fuel and time management tasks but fairly unanimous resistance to the same level of responsibility being delegated for the position management and EW management tasks.

Many of the aircrew voiced concern about the high integrity of the sensor information required to enable the decision processing used by the MMA to be totally valid. This is coloured to a certain extent by their current experience of RWRs but clearly the technology will need to provide consistent credible data in order to achieve aircrew acceptance.

It should be emphasised that the results contained in this paper were from a hands-off preview of an integrated MMA within a canned scenario.and it was not possible to examine performance measures associated with the allocation of MMA functions within a realistic, uncertain operational environment.

The MMA was still in a fairly embryonic form and the concept of task modelling was at an early stage but the' man machine team' approach certainly showed some promise as a method of developing an intelligent machine component.

FIG1    PROOF OF CONCEPT RIG
CONFIGURATION

# FIG2   PILOT AWARENESS AND MMA
## STATUS DISPLAY

**FIG3** **STATUS DISPLAYS**

# Achieving the Associate Relationship; Lessons Learned from 10 Years of Research and Design

Christopher A. Miller and Victor Riley
Honeywell Technology Center
3660 Technology Dr.
Minneapolis, MN 55418 U.S.A.

## 1. SUMMARY

The "associate relationship" is characterized by a mixed-initiative approach to *collaborative* problem solving between one or more human actors and a subordinate but semi-autonomous computer system with sufficient depth and range of intelligence and capabilities to encompass a full task domain. The associate relationship was first realized in a working system by the USAF Pilot's Associate (PA) programs. For nearly ten years, Honeywell has been leveraging lessons learned in the PA programs to extend the utility, effectiveness and achievability of the associate relationship. In this paper, we describe our work on seven recent programs which are extending the associate relationship by developing machine learning tools for associate knowledge base development or revision, by pioneering methods for reusing associate knowledge bases, by pushing associate applications in novel domains and for increasingly large numbers of users, and by performing basic research to understand when and how humans will decide to use an associate.

## 2. INTRODUCTION

The completion of the U.S. Air Force's Pilot's Associate (PA) programs in 1991 and 1992 resulted in, for the first time, the concrete existence of a new form of human-computer interaction: the associate system. The "associate relationship" is characterized by a mixed-initiative approach to *collaborative* problem solving between one or more human actors and a subordinate but semi-autonomous computer system with sufficient depth and range of intelligence and capabilities to encompass a full task domain (1). The associate relationship had been hypothesized in fiction for at least 40 years (Asimov's Robot books) and in serious interface design literature for at least ten years prior to the beginning of the PA programs, but the PA demonstrations were, we claim, the first time such systems had ever been seen in action.

As a result of the PA programs, we now know how to build one kind of associate system. What remains, therefore, is to expand and refine that knowledge; to learn how to make associate systems usable and useful in an expanded variety of real world settings.

Since the later phases of the PA programs, much of the research and development work on associate technology has taken that emphasis. The U.S. Navy's work on principles of adaptive aiding (2), Wright Laboratories continuing PA research, NASA's investigations of adaptive information management in the commercial flight domain (e.g., 3), and Onken's work on assistant systems for automobile drivers and general aviation pilots (e.g., 4) are examples of this trend, from among many others.

Much of our work over the past ten years also endeavors to improve the variety and utility of the associate relationship by increasing the breadth of domains for which an associate can be constructed, increasing the number of users it interacts with, developing tools to facilitate associate construction, or developing novel ways in which to leverage the utility of associate systems once constructed. The remainder of this paper will present brief accounts of seven associate technology programs at Honeywell, each of which leverages, and yet pushes the boundaries of, the USAF's Pilot's Associate programs. At the end of the paper, we draw general conclusions from this work for the future development of associate systems.

## 3. AUTOMATING KNOWLEDGE ACQUISITION FOR ASSOCIATES

### 3.1 Problem

By the mid-point of the PA programs it was clear that the quantity of knowledge required for a full associate system would be very large. Knowledge engineering was perhaps the largest problem for developing and, especially, for fielding an associate. Among the reasons for this were (1) the fact that the distributed, modular nature of the associate architecture made it necessary to have multiple, yet tightly coordinated, representations of the same knowledge (e.g., separate plan representations for a Tactics Planner, for a Pilot-Vehicle Interface, etc.), (2) the fact that PA knowledge engineering required multiple knowledge engineers working collaboratively, and (3) that the domain knowledge for a PA necessarily grows and changes over time as new tactics and equipment alter the utility of plans. Each of these factors contributes to the need for automated tools to support the initial acquisition of PA knowledge and its revision and maintenance over time.

### 3.2 Program

The Learning Systems for Pilot Aiding (LSPA) program was a 3.5 year effort, sponsored by Wright Laboratories, to address these questions by developing machine learning techniques to facilitate the acquisition or revision of the knowledge bases of an associate system. We used pilot behavior in simulation to provide learning examples: instances where a known goal was accomplished, but no plan known to the PA was used to accomplish it. We fed mission data captured by the simulator into an Explanation-Based Learning approach (5) to "explain" how pilot actions interacted with world states to accomplish the goal. This explanation was then generalized for use in future, similar instances, and then transformed, via a "smart translator," into the appropriate syntax for the PA Tactics Planner. The same explanation also served as input to a secondary "smart translator" which used a rule-based approach

(including knowledge about pilot-level actions for plan performance and the relationship between pilot-behavior and the way in which information is needed to support that behavior) to reason about the set of pilot information requirements for performing the new plan. The output of this second translator was an information requirements knowledge structure which was input into the Pilot-Vehicle Interface and used in the PA's information management reasoning whenever the new plan was activated.

## 3.3 Results

The LSPA program successfully learned both novel Tactics Planner plans and their associated Pilot-Vehicle Interface information requirements lists by observing pilot behavior. We developed knowledge sufficient to flesh out a branch of PA behavior (primarily concerned with defeating SAM sites) which had not previously been incorporated into the PA. Our approach reduced the time for developing novel tactical plans from approximately one week to about an hour, and human raters consistently judged LSPA's information requirements lists to be as accurate and more complete and consistent than those produced by human knowledge engineers.

## 3.4 Lessons Learned

LSPA showed that learning from pilot actions is entirely possible. Furthermore, the representation of causal knowledge in LSPA's learning system made it possible to more directly control, interact with and refine the knowledge bases of an associate. The development of explicit, causal knowledge in hierarchical form, along with smart translators for transforming changes in the causal knowledge base into changes in all other run-time knowledge bases, provides a good basis for exploring changes to the behavior of the associate. Knowledge can be input or modified directly at this level, and the higher-level, abstracted quality of the knowledge provides a good basis for mission rehearsal, training and other types of knowledge reuse.

The development of the translators for a LSPA-like system requires formalization of many aspects of PA design which were otherwise haphazard or intuitive. This formalization facilitates design accuracy and consistency and, to a larger degree, the ability to inspect and modify the design by later knowledge engineers. It can also be used to "deepen" explanations for why the PA behaves as it does since it accounts for why the knowledge by which the PA operates is the way it is.

Finally, the availability of such an automated knowledge engineering tool would greatly increase the completeness and flexibility of a fielded PA by enabling rapid modifications and additions to the knowledge bases upon which the PA operates. Our belief in the need for this capability has only grown in recent years as we have learned more about the individual differences and preferences between pilots and other operators, and about the standard operating procedures which differ from place to place and team to team within virtually all domains. It is becoming increasingly clear that it will be impossible to field a "generalized" associate with substantial aiding capabilities. Differences between individual operators are too great. Any fielded associate sys-

tem will need to be easily modifiable if it is to be accepted and useful.

While LSPA is a partial solution to the knowledge base development and maintenance problem for associate systems, more work is needed. Tools for the development, visualization, exploration, and rapid modification of associate knowledge are important and are under development in many places. Perhaps more important for the fielding of associates is the development of configuration and truth maintenance-type tools which will be capable of inspecting the knowledge bases of the associate and determining where discrepancies and inefficiencies exist, and then of suggesting methods for repairing them.

*For more information about the LSPA program, see (6), and its references.*

# 4. REUSING ASSOCIATE KNOWLEDGE BASES

## 4.1 Problem

The cost of acquiring, representing and revising knowledge for an associate is extremely large. Even though the use of a LSPA-like system holds promise for reducing that cost, knowledge acquisition for such a system is non-trivial in its own right. Methods for either increasing the speed and accuracy of knowledge acquisition or for leveraging associate knowledge once captured would both serve to alleviate this problem.

## 4.2 Programs

We have pursued both of these approaches in two recent programs. First, in a one-year follow-on to the baseline LSPA effort, we investigated the opportunities for acquiring LSPA domain knowledge directly from models (in the VHSIC Hardware Description Language--VHDL) of the structure and behavior of the electronic components of an aircraft. Work on this program centered on identifying the types of LSPA-relevant knowledge which were available in VHDL models, on developing methods for acquiring that knowledge from VHDL models and incorporating it into the causal knowledge of the LSPA domain theory, and finally, on attempting to quantify the time savings and accuracy increases available through such an approach.

A second, internally-funded effort has explored methods for using existing associate knowledge bases to leverage the design or redesign of future aircraft systems—specifically, the design of display formats. Since PA knowledge bases encode the mission-level goals and activities of the aircraft, along with specific behaviors required of aircraft subsystems to achieve the goals, they provide a ready resource for generating and evaluating system design requirements. In this effort, we examined the potential to optimize the set of plan-based information requirements and available display formats, represented in quantitative formats in the knowledge bases of the Pilot-Vehicle Interface, by applying a matrix analysis technique (Singular Value Decomposition) to determine the utility of each format and the information "coverage" of each task. A similar approach could be taken to rapidly analyze the impact of proposed novel formats.

### 4.3 Results

In the LSPA follow-on program, we successfully demonstrated the ability to acquire a variety of cockpit display-related information (e.g., latency values, control channels, pixel resolution, etc.) directly from VHDL models of the display hardware and to incorporate it into relevant portions of the information requirements logic of LSPA. Perhaps more importantly, we developed knowledge representations, based on Functional Reasoning (7), for linking the very low-level, hardware-based representations of the VHDL equipment models with the very high-level, largely human intentional state models used by LSPA and the PA modules. This, in turn, gave us the ability to reason in two directions: about the impact of equipment changes on the goals and plans of the human-aircraft system, and about the impact of the goals and plans intended for the aircraft on the low level hardware capabilities needed to enable those goals.

In the display format optimization program, we demonstrated the ability to analyze a matrix of display formats and their information requirements via the singular value decomposition approach-- and to provide useful, detailed, comprehensive feedback to format designers in minutes rather than the hours or days such feedback requires when done by hand. This approach proved particularly effective in providing data about the *overall* set of display formats and plans—a problem too large to be handled effectively by human engineers without such aids.

### 4.4 Lessons Learned

In the LSPA follow on effort, we learned that to reason about the capabilities of the low-level systems on the aircraft implied a dramatic broadening of the scope of knowledge needed for an associate. Instead of needing to know only the rough range of a FLIR, we now needed to know how a FLIR works at a very low level and how a host of world conditions interact to determine its range. While this is more than a PA needs to know, it is nevertheless exactly the set of concerns that must be taken into account in designing systems in the first place. Thus, while this approach is more than is necessary for knowledge acquisition for LSPA and PA, it points the way toward future, integrated design environments in which the design of aircraft hardware, of plans and goals for the aircraft itself, and of intelligent aids to assist in achieving those goals could go on simultaneously and leverage off of each other.

A common design environment, utilizing principles of knowledge sharing (cf. 8) should sew all these threads together, facilitating knowledge transfer in *both* directions: a proposed new system should draw its requirements from associate models of vehicle missions and functionality, reasoning about the impact of a new system on the vehicle should be done in conjunction with the knowledge resident in the associate, and finally, the capabilities and requirements of the new system should be incorporated directly from the system model into the knowledge bases of the associate, updating appropriate plans and actions as it goes.

Finally, this program has shown us that while the quantity of knowledge required for a full associate is extensive and difficult to acquire, it is also a valuable commodity in its own right. Our work on display format optimization techniques shows that existing PA knowledge bases can provide useful data for the design or redesign of display formats. The plan- and goal-related knowledge of the associate should be useful for the development and evaluation of other aircraft equipment, and for training, tactics development and evaluation, etc.

*For more information about the LSPA follow on, see (9), and its references. For more information about the Singular Value Decomposition approach to display optimization, see (10), and its references.*

## 5. NEXT-GENERATION ASSOCIATE

The largest associate system development effort underway in the U.S. today is the U.S. Army's Rotorcraft Pilot's Associate (RPA). RPA's goal is to develop, extensively evaluate and field test an associate system for use in a two crewmember advanced attack/scout helicopter. We are participating in the development of the Cockpit Information Manager with emphasis on the design of Information Management, Task Allocation and Taskload Estimation approaches for two crewmembers. Below, we single out two aspects of the RPA program which are extending our understanding of associate design and leveraging, extending, or applying lessons learned from prior associate programs.

### 5.1 Less Sequential Domain

The helicopter domain is generally less sequential than the fighter aircraft domain. The ability to slow, stop, and hover vastly increases the range of options about when and in what order tasks must be performed. Furthermore, helicopter operations tend to be less "scripted" and more susceptible to radical redirects than are fighter operations—and this is likely to increase as battlefields becomes more digitized.

One manifestation of this difference in initial RPA design is an increased emphasis on the representation of crew goals and intents, rather than straightforward scripting of tasks. Representations for goals and their importance stem from lessons learned in the PA programs (e.g., 11). Three reasons for this, and their implications for associate design, are described below.

First, the decreased sequentiality of behavior in the helicopter domain may make it more difficult for the associate to accurately track pilot intent and plans—and, therefore, more difficult to provide context-sensitive aiding. Since goals precede plans, both causally and in terms of generality, an explicit tracking of crew goals should be easier than knowing the particular plan being enacted. Knowing which goals are active will enable some accurate, goal-based aiding even when active plans cannot be determined.

Second, decreased sequentiality may make it more important for the crew to be able to communicate their intent directly to the associate. The associate should never *require* this input, but the benefit may justify minor increases in crew workload. In these

circumstances, we believe that a goal-based vocabulary of "intents" will provide the best mix of domain-validity, pilot-familiarity and usability by the associate. This will enable the crew to declare their goals and to set goals for the associate, below which it is capable and authorized to determine and enact courses of action.

Third, although many plans and goals may be "enabled" at a time (meaning that conditions are satisfied for them to be enacted), human resource constraints will continue to limit the number of plans which are actively being worked on. This distinction may be used to control information overload. Leveraging lessons learned in LSPA program, we will to support different types of information requirements for tasks which are enabled, versus those which are active. For example, an enabled but not active plan implies the need for information about success or failure conditions of the plan and about its temporal criticality. An active plan, on the other hand, implies the need for information pertinent to plan execution.

## 5.2 Aiding for Two Crewmembers
RPA is not the first associate program to examine aiding for multiple crewmembers, but it is certainly the most extensive. The fact that RPA's two crewmembers can and do trade tasks and responsibilities makes for both problems and opportunities. Problems stem from increasing the complexity of intent inferencing, information management, error detection, and adaptive aiding due to increased difficulty of knowing which task is being performed by which crewmember. Opportunities stem from the ability to use the associate to *facilitate* crew coordination through tracking the progress of task performance and through dynamically allocating some tasks between crewmembers to maintain acceptable taskloading and improve performance. When sequential dependencies exist between tasks being performed by separate crewmembers, the associate should be able to track task performance by both crewmembers and notify or aid them when the dependencies are in danger of being violated.

An additional question arises for the dual-crew context: how should an associate behave in the hierarchy of "command" for multiple crewmembers? In the single-pilot fighter domain, it was clear that the associate relationship required subordination to the human and consideration of that human's goals as paramount. In the multi-operator domain, however, in order to aid one crewmember it may be necessary for the associate to "command" the second and to prioritize one crewmember's goals over the others. It may be necessary to redefine the associate relationship as subordinate to *each* operator *within his or her sphere of authorization*. The associate will serve each crewmember only to the extent that it can do so without violating the plans of the other. Where conflicts exist, the associate has a duty to behave according to the established chain of command. One initial implication of this issue is the need for the associate to explicitly distinguish between when it is relaying (or acting on) a command from a superior crewmember to a subordinate versus when it is making a suggestion of its own to the subordinate.

## 6. ASSOCIATES FOR MANY OPERATORS

### 6.1 Problem
While RPA is developing an associate for 2 crewmembers with a common goal and command structure, the associate concept clearly holds potential for larger groups of people, whether they are engaged in a common goal or not. There are significant obstacles to overcome in defining associates to support multiple, potentially conflicting human actors. Geddes (12) has shown that plan-goal architectures developed for the single-operator PA can be extended to perform intent interpretation for multiple, conflicting human actors, but significant work remains to be done to define the overall behavior of an associate in multiple actor environments. Among the open issues are the role of the associate in a chain of command, the behavior of an associate for multiple actors with shared goals and beliefs, and the behavior of associate(s) for multiple actors with conflicting goals and beliefs.

### 6.2 Programs
We are beginning work on two programs which define associate behavior for large groups of human actors. The first is an associate-like aid to the detection and management of abnormal situations in large-scale industrial processing plants such as oil refineries or paper mills. In this domain, the activities of 30 or more human operators, distributed over a few square miles, must be coordinated to achieve the common goal of maximizing production while maintaining safety. While these individuals share a common goal, chain of command, and general beliefs about appropriate procedures and conflict resolution strategies, specific resolution strategies must be chosen and coordinated over large areas where the simple location and movement of individuals and resources represents a significant challenge.

The second program involves an associate-like system to facilitate information access and management for *all* of the stakeholders in a medical domain—doctors, nurses, technicians, clerks, staff, patients. etc. The breadth of this domain is such that the assumption of a common goal, beliefs or an agreed-upon chain of command is no longer tenable. Resource conflicts are inevitable, but it is doubtful that an associate could implement a resolution strategy which would be acceptable to all parties. Furthermore, the number, range, and distribution of potential users of such a system make it unlikely that a single associate could store or access sufficient coordinated information to meet all needs.

### 6.3 Issues and Approaches
For the industrial processing plant application, initial studies indicate that a more or less standard associate system approach will be feasible because all of the human actors share a common goal and beliefs about the chain of command and the appropriate plans to consider in various circumstances. The associate will be designed to provide specific but subordinate assistance for *all* of the human agents within their area of responsibility, but in its capacity of serving all, it may be required to issue instructions and commands to some. For example, in fulfilling the associate role for the plant supervisor, it

may well be required to give orders (with the supervisor's authority) to underlings.

For the medical information domain the situation differs, however. Here, stakeholders share no common goal or set of beliefs about the chain of command or the appropriate method for resolving problems. It seems untenable in this domain for a single associate to fulfill the associate role for all of the stakeholders. Instead, we are proposing to develop an architecture of a large number of associate-like systems, each implemented as a well-defined society of interacting agents and each serving an individual, a tightly coordinated group of individuals, or perhaps a general hospital function (such as an emergency room). Each individual's 'Interaction Society™" will fill the associate role individually for him or her, but among the society's duties will be the negotiation of priorities, duties, and resource access with other agent societies in the medical domain. Not all negotiations will go smoothly, and some problems will have to be passed to the human operator (or to even higher authorities) for arbitration, but these should be few in number. General hospital goals may be looked after by a hospital-level agent society, but of course, these goals will not be the same as the goals of any given doctor, technician or patient. We believe that our interacting society of associate-like agents is the only way to provide systems which fulfill the associate role for all of the participants in a domain this broad and heterogeneous.

### 6.4 Implications and Lessons Learned
The strongest implication for adapting associate systems to interact with larger groups of human actors is that the associate relationship may need to be subtly redefined. It appears impossible for the associate to be subservient to all in the same way that it is subservient to a single operator. In organizations where there are shared beliefs about goals, the chain of command, and appropriate methods, it may be possible for a single associate to play an omnipresent role, interacting with all humans, subservient to each within his or her purview, but capable of relaying commands and instructions to some in its capacity as an associate for their superiors. In less well-defined and homogenous domains, it may be necessary to abandon the concept of a single associate which maintains an associate relationship with all human actors, and instead emphasize the development of multiple associates (perhaps one for each human actor) which maintain a full associate relationship with their human counterparts, but which must negotiate relationships and resolve conflicts with other humans (and their associates) dynamically as they evolve.

## 7. WILL AN ASSOCIATE BE USED?

### 7.1 Problem
No matter how capable an associate is, its development will be wasted if the intended users choose not to use it. An associate system is a type of automation, and because no associate has received extensive enough usage to serve as a test case, we must look to operators' use of other forms of automation to estimate how operators will use associates.

Although recent work has focused on user trust in automation (e.g., 13), there are many factors that may affect an operator's decision whether to use automation. Some of these are: workload, task complexity, risk, automation reliability, individual differences (biases due to attitudes toward automation, training, or experience), and fatigue.

The large number of factors and the potentially complex and dynamic interdependencies between them make it challenging to anticipate whether operators will use an automated system. Of greater importance, however, is the need to anticipate the conditions under which an operator may use the automation inappropriately, such as relying on it when it is unreliable, or defeating it when the operator is unreliable.

### 7.2 Program
We have carried out an extensive series of experiments to investigate the influences of all the factors cited above on operator decisions to use or not use automation. The experiments used a simple computer game with two tasks that the operator had to perform simultaneously. One task could be automated at the subject's discretion, and the design of the tasks permitted independent manipulations of automation reliability over time, subject workload, task complexity, and the consequences of errors, to estimate the effects of risk on automation use. Measurements included levels of automation use across and within subjects, subject attitudes toward automation, subject confidence in own ability, actual subject accuracy at manual performance, and subject risk taking propensity.

Two of the experiments used university students as subjects and two used commercial airline pilots to determine whether difference in training and experience with automation might bias automation use. In Experiment 1, workload, task uncertainty, and automation reliability were manipulated independently and in combination to estimate how much effect each factor had on the automation use of university students. In Experiment 2, subjects' prior knowledge about automation behaviors was manipulated to untangle the relative contributions of uncertainty about automation accuracy states and future behaviors (or trust) to automation use. Of particular interest was how trust in automation developed during the initial experience and how it was affected by apparent automation failures and recoveries. In Experiment 3, the procedure of Experiment 1 was repeated with commercial transport pilots to estimate how differences in training and experience with automation in real systems might affect use of the automation in the study. In Experiment 4, the possible consequences of errors were increased to investigate the effects of risk on pilot automation use. All subjects were offered a cash award for whomever posted the highest score in each condition in order to promote attentiveness and to provide something of value that might be at risk.

### 7.3 Results
All experiments showed that workload did not affect automation reliance. Although this seems counter-intuitive, it agrees with the work of other investigators (e.g., 14). All experiments also showed that

task complexity and automation reliability had highly significant effects, and that fatigue and learning also played prominent roles.

Experiment 2 successfully untangled the contributions of automation state uncertainty and trust in automation and found that both contributed to automation use decisions, but did so differently in the initial, failure, and recovery responses. Experiment 3 demonstrated that commercial transport pilots were strongly biased in favor of automation use, with over a third of the pilots continuing to use the automation throughout failure periods. This was in contrast to the students who favored manual control and quickly stopped using the automation when it failed. Experiment 4 showed that higher levels of risk inhibited automation reliance after the second automation failure.

One highly consistent feature of all the results was the wide range of individual differences and strategies. It appears that individuals employ relatively simple automation use strategies that are influenced by small numbers of factors, but different people are influenced by different factors.

### 7.4   Lessons Learned

The results of these studies suggest that an associate system will be used in different ways by different people, and that individual differences may lead to different tendencies of over- and under-reliance on the system. When an associate system has a fairly high level of responsibility or authority, inappropriate reliance on the system can have important consequences, and the effects of these behaviors on system safety must be explored thoroughly before the system is fielded.

*This work is described in detail in (15).*

### 8.   CONCLUSION

The past ten years have seen dramatic steps toward realizing the associate relationship. The USAF PA brought about a working associate system for a single operator. Our recent work is extending and leveraging associate research in a variety of directions. The development of associates in novel domains and for multiple operators (by ourselves and others) is encouraging in that it shows that a great deal of the PA concepts and architectures are relevant, but it is also teaching us about the similarities and differences as each new domain is tackled. Simultaneously, approaches are being developed to make associates easier and more profitable to build. While significant challenges remain on this front, it is encouraging to see that associate knowledge bases appear to have extensive utility outside of the associate system itself. Finally, as the development and fielding of associate systems proceeds, significant basic research remains to be done to establish when and how associates will be most useful—and most appropriately used.

### REFERENCES

1. Yadrick, R., Judge, C., and Riley, V. (1990). "Decision support and adaptive aiding for a battlefield interdiction mission", in the *Proceedings of The Human-Electronic Crew: Is the Team Maturing*, Ingolstadt, Germany.

2. Morrison, J. G., Cohen, D., and Gluckman, J. P. (1993) "Prospective principles and guidelines for the design of adaptively automated crewstations." *Proceedings of the Seventh International Symposium on Aviation Psychology*, Columbus, OH, pp. 172-177.

3. Shalin, V. Geddes, N. & Miksell, B. (1992). Information Management for the Commercial Aviation Flight Deck. Presented at the Aerospace Systems Technical Group Poster Session at the 36th Annual Meeting of the Human Factors Society, Atlanta, GA.

4. Kopf, M. and Onken, R. (1992). DAISY: A Knowledgeable Monitoring and Warning Aid for the Driver on German Motorways. In *Proceedings of the IFAC Man-Machine Symposium.*

5. Dejong, G. and Mooney, R. (1986). Explanation-Based Learning: An Alternative View. *Machine Learning, 1.* pp. 145-176.

6. Miller, C. and Levi, K. (1994). Linked-learning for Knowledge Acquisition: A Pilot's Associate Case Study. *Knowledge Acquisition, 6.* pp. 93-114.

7. Sembugamoorthy, V. and Chandrasekaran, B. (1986). Functional Representation of Devices and Compilation of Diagnostic Problem-solving Systems. In J. Kolodner and C. Riesbeck (Eds.), *Experience, Memory and Reasoning.* Erlbaum; Hillsdale, NJ. pp. 47-73.

8. Gruber, T., Tenenbaum, J. and Weber, J. (1992). Toward a Knowledge Medium for Collaborative Product Design. In *Proceedings of the 2nd International Conference on AI in Design.* Pitts-

burgh, June 22-25. Kluwer Academic Publishers, pp. 412-432.

9. Levi, K., Moberg, D., Miller, C., and Rose, F. (1993). Multilevel Causal Process Modeling: Bridging the Plan, Execution and Device-Implementation Gap. In *Proceedings of Applications of AI, XI*. Orlando, FL, April. SPIE. pp. 240-250.

10. Miller, C. and Morton, B. (1994). A Singular Value Decomposition Approach to Information Presentation Format Optimization. In *Proceedings of the 38th Annual Meeting of the Human Factors and Ergonomics Society*. Nashville, TN; October.

11. Sewell, D. and Geddes, N. (1990). A plan and goal based method for computer-human system design. *Human Computer Interaction: INTERACT-90*. New York: North-Holland. pp. 283-288.

12. Geddes, N. (in press). A Model for Intent Interpretation for Multiple Agents with Conflicts. To appear in *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*. October, San Antonio.

13. Muir, B. (1987). Trust between humans and machines, and the design of decision aids. *International Journal of Man-Machine Studies, 27*, 527-539.

14. Harris, W., Hancock, P., and Arthur, E. (1993). The effect of taskload projection on automation use, performance, and workload. In the *Proceedings of the Seventh International Symposium on Aviation Psychology*. Columbus, OH.

15. Riley, V. (1994). *Human Use of Automation*. Unpublished Doctoral Dissertation, University of Minnesota, Minneapolis, MN.

# Principles of Interaction for Intelligent Systems

John P. Zenyuh
Ronald L. Small
John M. Hammer
Alan D. Greenberg

Search Technology
4898 South Old Peachtree Road, #200
Norcross, Georgia 30071-4707

## 1. SUMMARY

Aircraft manufacturers, both commercial and military, are building automated systems of ever-increasing complexity and authority. There are several driving forces behind this trend: operating cost reductions, safety improvements, and performance optimizations. The increasing power of automation is redefining the role of the pilot and greatly complicating the pattern of interactions between the pilot and the aircraft systems.

Concern about enigmatic automation has been raised in both the operational and technical communities. Previous research has investigated the issue of operator trust in automated systems, identifying such key behaviors as predictability and repeatability. This paper addresses principles of interaction and their correlates, specific authorizations, as mechanisms for defining and implementing these expected behaviors. The process for deriving specific authorizations from global principles is one of iteratively incorporating contextual detail within the rules, achieving greater specificity with each successive pass. The design of a prototype Hazard Monitoring system is discussed as a specific application of this rule derivation process.

## 2. INTRODUCTION

The modern flight deck continues to become more dominated by automation of increasingly higher levels of authority; particularly in the area of flight management. Aircraft now operate more efficiently, but at the cost of distancing pilots from the flight control process. The flight crew spends a significant portion of its time and attention managing the automation rather than directly controlling the aircraft. As the complexity of these automated systems has increased, so too have the intricacies and interdependencies of their behavior.

Unfortunately, this continued expansion of automation throughout the flight deck has been accompanied by relatively little feedback about the specific behaviors and "intentions" of automation. In the absence of provisions for communicating their internal state, system behaviors inevitably diverge from the pilot's expectations, causing surprising and potentially dangerous behavior.[1] Each time the system surprises the pilot, he/she tries to understand how his/her model and the intentions derived from it differ from the actual states of the automation. This leads to the commonly expressed sentiment among pilots, "What is the system doing to me now?" Additionally, the actions themselves frequently have consequences that conflict with the pilot's intentions. The result is that high-authority, automated systems often exhibit enigmatic behavior that surprises the crew, undermines trust, increases system monitoring and mental workload, and decreases the margin of safety as pilots' own behaviors become correspondingly reactive.

Flight Management Systems (FMSs), descendants of early Inertial Navigation Systems (INSs) and auto-pilot couplings, are one of the most pervasive high-authority, automated systems in today's flight decks. The flight management system itself is divided into three functional areas: navigation, thrust/performance management, and guidance. Each system can operate independently or be coupled with one or both of the other systems to manage aircraft control. Figure 1 below shows these components and their connections.

The pilot can independently disengage any of the subsystems, leaving the automation only partially in control of the aircraft. Indication that the systems are coupled appears as simple enunciator lights on the primary flight display (PFD) and the absence or presence of certain information on the mode control panel (MCP). System engagement is direct; the pilot presses the appropriate MCP buttons. Disengagement, in contrast, is subtle; it occurs whenever the pilot manipulates an MCP control which overrides another FMS mode or

Figure 1. Typical flight management automation architecture.

when a previously set MCP control threshold is met. Additionally, the behavior of the system is dependent on the data previously entered by the flight crew. In general, the interrelations between the flight management, auto-thrust, and auto-flight systems are complex, subtle, and often poorly communicated. Such cryptic behavior from a human team member would be unacceptable.

## 3. BACKGROUND

Concerns that unpredictable automation will result in crew distrust, thus reducing any potential benefits, are driving a growing interest in the nature of high-reliability systems. Air traffic control towers and flight deck environments are working examples of complex systems or teams that have remarkably low rates of failure. In each case, the human teams have developed patterns of interaction that effectively keep the team informed about task allocation, state changes, and critical events without the need for verbose exchanges of information. The teams work to keep the task load even and manageable, the level of authority appropriate to the task, and the current team goal apparent to everyone.[2] The team interactions promote trust and learning by being both reliable and error tolerant.

Just as a crew member must have trust in the crew, a significant factor in the usability and acceptability of automation is the operator's trust in the system. Trust, as described by Muir, develops in three stages.[3] Initially, trust is based on predictability of individual behaviors. Secondly, trust is based on dependability, which can be thought of as a measure of the degree of system reliability in any situation. Finally, the user develops faith that the system will continue to be dependable in the future, possibly because of its

ability to improve itself or its robustness. As such, specific authorizations derived from the appropriate principles of interaction will foster trust in the system, which in the case of current cockpit automation, is sorely lacking.

While high operator trust is an important goal for successful system implementation, the quintessential objective of high authority automation is still to optimize operator-system task performance. In a study expanding upon Muir's original work, Lee and Moray investigated operator performance in response to various automated system behaviors.[4] They discovered that operator performance measures describe a remarkably similar curve to that of operator trust given equivalent behavior of an automated system over time. Specifically, both trust and performance exhibit gradual, but steady, increases as an operator interacts with a system that exhibits reliable, predictable behaviors. Additionally, both measures show marked drops in response to transient, anomalous system behaviors and a significantly suppressed 're-learning' curve in response to more permanent anomalous conditions. Therefore, in addition to fostering trust in the human-system interrelationship, automated functions based on consistent, well-defined, well-communicated principles of interaction should also result in better, more stable performance.

## 4. RULES OF INTERACTION

Flight crew members communicate in a well-defined format in order to share information, establish expectations (both commitments and responsibilities) for task performance, and negotiate the division of work. The goal of their task-oriented communication is to develop a shared situation awareness and to ensure predictable, consistent team behavior. If conflicts arise between the goals of the crew and the current intentions of the pilot (e.g., he tells them to navigate toward an unknown waypoint), they clearly state the problem and ask for confirmation. Creating an environment of shared understanding, trust, and commitment and then acting within it leads to efficient and highly reliable team performance. How a crew member decides if he or she is authorized to communicate or act is predicated in protocols that guide flight deck behavior. These protocols, and the mechanisms used to maintain them, can be embodied in simple rules of interaction. These rules can then be used to restructure and improve the pattern of interactions between humans and machines.

163

## 4.1 Communication Protocol

Our interest in this problem arose from our work on the Pilot's Associate (PA), an intelligent, high-authority system for aiding the pilot of a single-seat tactical aircraft. The purpose of the program was to explore how artificial intelligence might aid the pilot if given authority to automatically reconfigure displays, monitor for pilot error, propose plans of action, and execute virtually any action or procedure on his/her behalf.[5]

Our observations of, and conversations with, tactical pilots suggested that a relatively sophisticated model would be required to sufficiently identify conditions where PA was authorized to act. One form of evidence for this was the relatively infrequent communication (e.g., once every 10-15 seconds) between experienced crews in two-seat tactical aircraft. The minimal need for communication suggested that each crew member has some type of model that permits considerable action with minimal explicit communication. A similar kind of model and level of performance would be required to enable an electronic crew member to function as well as a human crew member.

One approach to solving this problem is to provide the intelligent automation with psychological models of the pilot. Work in psychology suggests that it should be possible to compute the pilot's intentions, workload, and focus of attention then use this information to make the automation's actions more comprehensible. However, our experience indicated that these models are not effective. First, they are based upon inaccessible information (i.e., mental and neural states of the pilot) and implicit intuition and judgment that the machine is, as yet, unable to provide. And second, pilot interviews indicated they would more readily accept the aiding if the decision were based on situational aspects that are more salient and observable.

The issue then, was to formulate a set of rules governing system interaction based on tangible measurements that can be ascertained concurrently with, or predicted in advance of, the exercise of authority.

## 4.2 Human-Human Rules

Initially, it would appear that a reasonable set of rules identified from crew interaction could serve as a basis for a similar set of rules for an intelligent interface. However, the intelligent interface is not human and, therefore, the rules used by humans may be only partially suitable for it. For example, pilots may never be willing to delegate as much authority to automation (however intelligent) as to other human crew members. Still, identification of the rules used by humans is an important first step to a set of rules for intelligent interfaces. These rules of interaction would govern some of the most visible and significant behavior of the intelligent system, especially that aspect to which crew members are most sensitive.

Our initial pilot interviews during the PA effort revealed an abundance of proposed principles for defining human-automation interaction. A representative sample follows.

- The human has ultimate control and can override the associate system at any time.

- The associate system must follow the lead of the human operator. The associate monitors and supports the human, not the other way around.

- The system must never seize control from the pilot.

- The operator should have the option of turning off part or all of the associate system. For example, error monitoring may be turned off when the operator is involved in innovative, unconventional behavior.

Significantly, a number of these principles were relatively global in nature, relying on abstract judgments and devoid of contextual information. For example, principles tended to refer to "high pilot mental workload," "stress" or "innovation." The result is a set of universally accepted global principles, but no specific, well-defined authorizations for automation behavior.

There are several solutions to this problem. One solution is simply to discard these rules. A second is to make the rules more specific. For example, the rule might incorporate notions of whether the pilot is entering an atypically dense volume of airspace or experiencing a particularly busy portion of the flight. Of course, such clauses are only an approximation of the condition in the original rule. A third solution is to *situate* the rule.[6] The idea here is that in a more specific situation, it is possible for a rule to be less ambiguous. For example, instead of "reduce workload," the system should use rules like "on final approach, landing gear should be fully extended." In this case, final approach conditions can be determined unambiguously without the

need for encoding intuition or judgment. Then, if the gear are not extended at the appropriate time, the pilot can be alerted to the unsafe condition, regardless of the reason for his/her inactivity. Therefore, even if the pilot's general state or intent is unknown, well-defined rules of authorization can be used to respond to specific situations of interest.

### 4.3 Rule Refinement

The process for deriving these situated authorizations is to iteratively incorporate greater contextual detail into the initial global principles. For example, consider the principle, "The system must never seize control from the pilot." While the spirit of this principle is well understood, by its abstract nature the definition of the principle is inherently incomplete. A more accurate definition would be, "The system must never seize flight control from the pilot, except to save the pilot's life." However, it too lacks sufficient situational context to support implementation. Continuing the example, "The system must never seize thrust and/or flight management control from the pilot unless ground impact is imminent and recovery is beyond the capability of the pilot." Additional iterations would be required to more clearly define the concepts of "imminent ground impact" and "pilot capability." The final step in this process is to rephrase the principle as a specific system authorization: "In the event of imminent ground impact, based on current trajectory and aircraft configuration, where recovery requires thrust and control surface inputs exceeding the physical and cognitive abilities of the pilot, the automation is authorized to institute recovery and maintain straight and level flight." The result is a domain-based rule which defines clear, predictable, and implementable automation behavior.

### 5. THE HAZARD MONITOR

A key technology, conceived during the initial Pilot's Associate effort (continued through the sponsorship of NASA Langley), that has taken advantage of our various lessons learned regarding principles and authorizations for automation behavior is the Hazard Monitor (HM). The primary goal of the HM module is to unobtrusively track the various activities of the flight crew and to notify them of any aberrant or unexpected behaviors before such actions result in negative consequences.[7,8]

Specific HM monitoring behaviors are guided by rules of interaction ascertained through knowledge acquisition with commercial pilots. An important consideration here is that HM does not employ a rule-based technology to perform its functions.

HM's behavioral protocol is embodied in a series of monitoring networks. Each network, in turn, is composed of related expectations of flight crew behavior. A well-defined set of initiators and terminators ensures that only the appropriate types of monitoring occur at any point during the flight. Expectations of pilot behavior and appropriate automation responses (i.e., continued monitoring, remedial directives, or control actions) are situated throughout the monitoring process by discrete temporal- and state-based events.

A detailed example of HM's implementation may help clarify this point. Consider the principle that HM is to aid the pilot in managing significant altitude changes. Altitude monitoring is to occur at all times during flight. "In-flight" is later defined as from weight-off-wheels (i.e., airborne) to weight-on-wheels (i.e., "on the ground" condition). These two criteria are incorporated as the initiator and terminator of the altitude monitoring network. Specifically, HM is to monitor for expected standard and local altimeter settings on climb and descent. Climb and descent are then more precisely defined as transition altitude (i.e., 18,000' in the USA and 6,000' in some parts of Europe). Since HM can vary its level of response (e.g., monitor, enhance, warn, etc.), pre- and post-threshold events (i.e., 500' above and below transition altitude) are chosen as opportunities for HM intervention. The result is a monitoring network that exhibits clear, consistent behavior based on a defined set of rules which govern effective flight deck interaction. Figure 2 shows the structure of an HM monitoring network for this principle.



Figure 2. Example Hazard Monitor Net.

## 6. CONCLUSIONS

The trend in modern flight deck design is clearly to increase automation and its authority in aircraft. With the cost of each transport aircraft now exceeding $100 million, the cost of litigation resulting from a crash, and the non-quantifiable cost of human life lost, the significance of the need to understand the relationship between pilots and automated systems cannot be overestimated. As such, the problem of appropriate automation is a concern to the scientific community. Bemoaning this irreversible trend is simply an inappropriate response. It is our responsibility to shape technology in useful ways. Some have advocated that more attention be given to human factors design of automation. While human factors goals apply to the functionality of automation, its methods and accumulated knowledge are more appropriate to displays and controls. What is needed is a complementary "human factors of functionality" that is concerned with the behavior of systems. This can be done best by developing specific guidance on appropriate behavior for automation. The process of identifying principles of interaction and the subsequent derivation and implementation of specific authorizations is an initial, yet definitive, step toward the creation of predictable, consistent and trustworthy automated systems.

## 7. ACKNOWLEDGMENTS

## 8. REFERENCES

1. Rouse, W.B., Cannon-Bowers, J.A., and Salas, E. "The Role of Mental Models in Team Performance in Complex Systems." IEEE Transactions on Systems, Man and Cybernetics, 22, 6, November/December 1992, 1296 -1308.
2. Selcon, S.J., and Taylor, R.M. "Psychological principles for human-electronic crew teamwork." Proceedings of the Second Joint GAF/RAF/USAF Workshop on Human-Electronic Crew Teamwork. Ingolstadt, Germany. September 1990, Paper 21.
3. Muir, B.M. "Trust between humans and machines, and the design of decision aids." International Journal of Man-Machine Studies, 1987, 27, 527-539.
4. Lee, J., and Moray, N. "Trust, control strategies and allocation of function in human-machine systems." Ergonomics, 1992, 35, 10, 1243-1270.
5. Hammer, J.M., and Small, R.L. (to appear). "An intelligent interface in an associate system." In Human/Technology Interaction in Complex Systems (W.B. Rouse, ed.) Greenwich, CT: JAI Press.
6. Suchman, L.A. "Plans and Situated Actions." Cambridge University Press: Cambridge, 1987.
7. Greenberg, A.D., Small, R.L., Zenyuh, J.P., and Skidmore, M.D. (to appear). "Monitoring for hazard in flight management systems." European Journal of Operations Research.
8. Small, R.L., Zenyuh, J.P., Skidmore, M.D., Greenberg, A.D., and Hammer, J.M. "A prototype flight management system error monitor." (Tech. Report STI-TR-9209-001). Norcross, GA: Search Technology, Inc., August 1993.

# The "Copilote Electronique" project :
## First lessons as exploratory development starts.

Thierry JOUBERT[1], Stéphane E. SALLÉ[2], Gilles CHAMPIGNEUX[3],
Jean Yves GRAU[4], Pierre SASSUS[5], Hervé LE DOEUFF[6].

1 : MATRA DEFENSE - DEA/SIA, 37 Avenue Louis Bréguet, BP 1,
    78146 VÉLIZY VILLACOUBLAY CEDEX France, Tél : 33 (1) 348884612, Fax : 33 (1) 348884455

2 : SAGEM - Centre d'Eragny sur Oise, Avenue du Gros Chêne, ERAGNY, B.P. 51,
    95612 CERGY PONTOISE CEDEX France, Tél :33 (1) 34.30.50.50, Fax :33 (1) 34.30.50.96

3 : DASSAULT AVIATION - DEA/DIA2/AIA, 78, quai Marcel Dassault, B.P. 300,
    92214 SAINT CLOUD France, Tél : 33 (1) 47.11.40.37, Fax : 33 (1) 47.11.52.83

4 : IMASSA, CEV Brétigny
    91228 Brétigny, France, Tél : 33 (1) 69883373, Fax : 33 (1) 60840448

5 : SEXTANT AVIONIQUE- DAC/DT/EF, Aérodrome de Villacoublay, BP 59,
    78141 Vélizy Villacoublay, France, Tél : 33 (1) 46297187, Fax : 33 (1) 46297766

6 : DASSAULT ELECTRONIQUE - LTI/LOC, parc d'activité de Pissaloup, Avenue Jean d'Alembert,
    78210 TRAPPES ELANCOURT, France, Tél : 33 (1) 34816252, Fax : 33 (1) 34817205

## 1. SUMMARY
In spite of great benefits expected, Knowledge Based System (KBS) approaches are not so easy to apply in the pilot assistance field. This paper presents first lessons learnt as a new phase of the French project "Copilote Electronique" has started in 1994, under the form of an exploratory development, advantages and drawbacks of existing methodologies have been compared. As a result, emphasis has been put on ergonomics design rules and on a project life cycle adaptation aiming at insuring better responses to pilots demands and fears ... and so preparing, we hope, successful operational evaluations.

Key words : Pilot Assistance, Knowledge Based Systems, Software Engineering, Life Cycle, "Copilote Electronique", Electronic Crew Member, Ergonomics Rules.

## 2. INTRODUCTION

This paper will present actual status of an Electronic Crew Member System project, called in French "COPILOTE ELECTRONIQUE".

Chapter 2 describes general goals and characteristics of the project "COPILOTE ELECTRONIQUE".

Chapter 3 first recalls design difficulties encountered in such projects : symbolic programming is not magic, knowledge acquisition is not so easy, process automation has to keep user or pilot in the loop, internal and external co-ordination problems arise and specific ergonomics rules are to be applied.

Chapter 4 proposes a discussion on rapid prototyping life-cycle model; its advantages and known drawbacks and possible solution from classical software engineering field. The knowledge acquisition and software engineering life cycle chosen for "COPILOTE ELECTRONIQUE" is illustrated by the inverse V-model.

Chapter 5 summarises first lessons learnt during the preparation of the project launching and the first six month of the new development phase.

## 3. THE "COPILOTE ELECTRONIQUE" PROJECT
[Air & Cosmos]

System and software design methods that have been used for current generation fighters are facing more and more difficult challenges and may encounter their own limits within ten or twenty years from now.

Various embedded functions, such as navigation, piloting, aircraft status, weapons system management, and in some extension sensors management have been successfully automated by classical software engineering methods, but the addition of such separate and independent automated functions become more and more difficult to control in real time situations by human pilots.

Nevertheless, as critical decisions are to be taken on uncertain or tactical aspects of mission, aircraft designers often rely, and this is currently required by Air Forces, on pilots judgement.

But as automated functions are intended to increase in number and *complexity*, in the foreseen tactical context tactical characterised by a great number of various possible threats, with electronic war systems and new sophisticated weapons, we may think that future pilots will have some difficulties with this combinatory explosion of information sources unless being assisted in their reasoning tasks.

An expert assistance system, as the "COPILOTE ELECTRONIQUE" will have to absorb high rates of raw information, select and highlight the more crucial ones, before initiating dialogue, in a manner adapted to current situation and mental load of the pilot. As we say, such a system should only present *pertinent information* and offers a restricted actions choice to the pilot, on which, after selection by the pilot, it will have to examine all consequences before execution.

The "COPILOTE ELECTRONIQUE", initialised in 1986 by the French DGA ( "Délégation Générale de l'Armement"), aims at introducing, within a 2010 horizon, expert or knowledge based systems in combat aircrafts. Far from replacing human pilots in the cockpits, such an

sophisticated electronic crew member should be considered as a very *high level dialogue function* between man and machine.

In fact, this project took a new acceleration in 1994 spring, when the Technical Service for Aeronautical Telecommunication and Equipments of French DGA (STTE) decided the funding of an exploratory development for RAFALE standard SU2, which is the Rafale standard that will enter French Air Forces in 2004. SU2 standard will benefits of all radar modes, counter measures and a front infrared detection and tracking system; besides pilot will use helmet mounted displays.

The goal of this new phase, launched for a three years duration, is a ground simulation, without hard real time constraints, to demonstrate the "COPILOTE ELECTRONIQUE" of Rafale SU2 in situation of strike and escort missions, with low altitude penetration constraints.

## 4. DIFFICULTIES FOR PILOT ASSISTANCE SYSTEM DESIGN

The complexity of problems in the aeronautic field led some designers to propose KBS methods as an easier and more comfortable programming solution than "classical laborious and error prone" software development. With those symbolic software, developers quickly produced promising prototypes in view of production systems

delivery. But as no specific software engineering methodologies were generally applied, it became obvious that desirable high quality and maintainable systems were not reachable. Since then a relative disillusion is felt in the aeronautic community.

To understand KBS interest and complexity of development, one has to take into account the importance of human expertise in the design process. Human experts need to be considered as full members of the project team involved at each stage of the development : early description of the problem domain, requirements definition, design, description of performed tasks, ideas of new man-machine dialogue, validation, end-use,... This central role influences the developing environment and suggests the modification of the infrastructure. We can also mention that KBS pretends to level up the knowledge representations so that the human specialists can understand what is in the machine. It requires new concepts like objects, plans, heuristics, agents... The expert may even want to flip from one representation to another. So that, knowledge engineers have to walk their way through a very large set of representation schemes Pilot Assistance systems must present specific characteristics : they must be real-time systems (involving most of the time some temporal reasoning), embedded on-board aircrafts (satisfying CPU, memory, ... restrictions), most frequently multi-expert, deeply integrated into their environment and keeping the end-user in-the-loop.

In [*AGARD 1993*], a functional analysis of Knowledge Based Systems in Guidance and Control (G+C) field is stated



Figure 1 : Guidance and control system functional decomposition

168

(figure 1) and seems quite adapted to pilot's assistance field. In particular, the classification in three main engineering domain is quite pertinent : situation assessment (monitoring, diagnosis), plan management (plan selection and execution), co-ordination (external and internal) and central node plan generation and is a good summary of problems encountered.

A more detailed discussion on specific difficulties for intelligent assistance design in G+C field can be found in [Sallé1993 ].

Non respect of ergonomics rules is the most current explanation for KBS applications failures, in the field of industrial processing assistance. So IMASSA (Centre for Medical Studies and Research in Aerospace) was in charge to define "user oriented rules" that has to be used from the design phase [Amalberti et al 1990 ].

Those rules can be summarised as follows :
    • pilot anticipates and needs anticipation assistance on contrary of "classical engineer designed" assistance which are often too reactive,
    • pilot's decisions reflect often compromises between mental load and ideal response to the situation, so pure optimality is not to be researched if pilot has no sufficient time to understand,
    • following their own personal skills, different pilots may organise work differently, assistance must be adapted to these skills,
    • assistance must be homogeneous, and it will be preferable to rely on specialised expert for each operational domain (e.g. STRIKE or AIR DEFENSE expertise) so resulting assistance will produce constant understanding interpretation model that will avoid surprises for pilot
    • assistance must know and respect its own limits
    • system design may use "what if" approach to be less reactive
    • dialogue must be adapted to context, pilot intents and pilot load
    • dialogue must be space oriented and interactive, better use vocal media than written, but avoid saturation,
    • respect logic of pilot understanding, that means rely on the understanding model designed with expert pilots.

## 5. DISCUSSION ABOUT SOFTWARE LIFE CYCLE IN KBS METHODS

The use of rapid prototyping to build a complete system, using one of the many software packages available on the market, has been a frequent technique in the early studies of "COPILOTE ELECTRONIQUE" project. This method also known as "evolutionary prototyping" or "iterating prototyping" was deduced from experimental approaches to KBS development.

It consists in iterating the cycle :
    • knowledge acquisition,
    • implementation : knowledge modelling & coding,
    • validation,
    • test with the expert(s), until there is no more knowledge to capture.

Detailed discussion on obvious advantages but also on subtle drawbacks of this life-cycle methodology can be found in [Sallé1993 ]. In summary, let's say that this rapid prototyping methodology is not relevant for complex and embedded systems for which a good architectural design is needed.

As a consequence of risks encountered in final system integration phase, when rapid prototyping method is followed, alternate methods were looked for the project.

Some theoretical studies tend to consider the KBS as an ordinary software production problem with its overall analysis prior to any implementation. KADS is the leader of this new way of design [Hickman et al 1989], MOISE [Ermine 1992] is another example.

Based on a model driven-approach, these methodologies have much in common with conventional software development methodologies : they prescribe phases, stages and activities, models, documents and deliverables (see figure 2).

To maintain the benefits of dialogue with experts, a KBS approach iterating design illustrated by the spiral model (figure 3) was introduced [Hickman et al 1989]. This model may be considered as a overall life-cycle model and may be inserted at the top of the V-model depicted in figure 2. We



Figure 2 : COPILOTE ELECTRONIQUE V-cycle

use this spiral model to show all efforts made in the early phase of "Copilote Electronique", before actual project launching (see chapter 5.)

Current knowledge engineering practices heavily depends on interview techniques and the collection and analysis of notes. The process, although valuable, is slow and frequently paces the development activity.

There is general agreement that considerable gains in speed and efficiency can be achieved by improving both tools and methods. But additional efforts, for software designers, in knowledge acquisition and knowledge elicitation from expert pilots are not to be minimised, we think, as it is the *main benefit of KBS oriented pilot assistance design*. This can be illustrated by the climbing branch of figure 2 inverse V-model (following some ideas from [*Dieng 1990*]), this climbing branch being more important than in classical V-cycle, where existence of a clear specification is always the starting point of process and future disappointments.

# 6.LESSONS LEARNT IN "COPILOTE ELECTRONIQUE" PROJECT

At the phase of problem understanding and needs definition of the "Copilote Electronique" program (see figure 3), three main risks were identified :
- Is it possible to capture enough expertise to create a real assistance for pilot reasoning ?
- Is the KBS technology mature enough for a real development ?
- Is the French Aerospace community able to integrate such a new concept in a current avionics system design ?

A survey of existing efforts (national as well as international) was initiated to give answers to these questions. An extensive work sponsored by DRET (French Defence Advanced Research Agency) and realised by IMASSA (Centre for Medical Studies and Research in Aerospace), with Mirage F1 Recce pilots, gave a lot of clues

resulting in the identification of various pilots behaviour correlated with pilots profile. It brought to the front scene the pilot's "meta-knowledge" (specialised technical education) which influences mission planning as well as reflex reactions. A joint work with Dassault allowed to map pilots intuitive aspiration for new assistance with the reality of Mirage 2000 and Rafale advanced design.

In the five companies participating to the project, team was selected among people having experimented KBS methods, various references can be found in following list : [*Aubry et al 1988* ], [*Brunessaux et al 1992* ], [*Champigneux et al 1989* ], [*Desard et al 1991* ], [*Joubert et al 1990* ], [*Lementec et al1991* ], [*Morillon et al 1988* ], [*Salle et al 1989* ]. All these studies gave a strong background to investigate the cognitive aspects related to the various areas of expertise associated with decision aid, and the computer science aspects relative to the implementation of artificial intelligence techniques in real time airborne systems. They were studied to better plan the future development of the "Copilote Electronique".

Another task was to allocate assistance modules development between different companies (for the modular architecture see the "pyramid" figure 5) .

This allocation process happened to be difficult and risky :
- difficult because the concept being still fuzzy and. each company having a large scope of competencies the limit of each assistance domain were controversial ;
- risky because each partner had its own goals, culture, and methods and the system design could have resulted in a collection of nice but incoherent functionalities.

The rule chosen was that the company must possess high motivation for the specific assistance field chosen, but this must not be considered as a definitive allocation for future industrialisation.



Figure 3 : Spiral Model applied to Copilote Electronique

170

**Figure 4 :** Copilote Electronique architecture: the pyramid

| | | |
|---|---|---|
| •System Status Assessment and Management | -> | SAGEM |
| •Mission Conditions Assessment and Mission Management | -> | SEXTANT AVIONIQUE |
| •Tactical Situation Assessment and Management (Air threat and offensive Weapons) | -> | MATRA DEFENSE |
| •Tactical Situation Assessment and Management (Ground threat and defensive Counter Measures) | -> | DASSAULT ELECTRONIQUE |
| •Pilot Assessment, action plans assessment, relevant information management and man machine interface | -> | DASSAULT AVIATION |
| • Ergonomics rules and knowledge acquisition methods and verification tasks | -> | IMASSA |

**Figure 5 : Final allocation is the following :**

Several actions were initiated to reduce foreseen difficulties. An international co-operation was carried by Dassault with Lockheed in order to confront the French approach with the US team experience and avoid traps experimented in previous experiments [Smith et al 1988 ], [ Rouse 1991 ]. Our experience in the development of the Copilote Electronique as well as the technical exchanges with Lockheed Pilot's Associate team, led to the fact that conventional knowledge engineering techniques using questionnaires and interviews are not sufficient *to provide implementable and secured knowledge for Pilot aids. The* investigated knowledge engineering techniques, in order to tackle the problem of such a complex AI system development, should be used for expertise initial design, then supplemented by extensive knowledge evaluation and correction in simulator. With IMASSA, a *specific method for eliciting and formalising pilot's expert knowledge* was studied and is used.

Another very important technical issue is the definition of a *common "plans and goals" exchange language* between

all specific assistance modules, and great efforts are to be made to maintain this common message glossary.

Finally, a unifying technical principle was adopted to facilitate the architecture design via the *intent planning paradigm.* This principle is essential to fulfil general ergonomics constraints of chapter 3.3 : assistance must not participate to the signalled existing overloading factors. Intent recognition is a challenging but promising direction and can be made easier by extended preparation mission plans and procedures (for each pilot activity) that will be perhaps the new "automated and personalised" check lists version of the future.

## 8. CONCLUSION[1]

The technology is available today to provide viable knowledge system solutions to well-chosen and well defined problems. It can be expected to see more and more successful projects on such on-board applications, as both the research, the technology and engineering skills of application developers improve.

But this process may be slower than was though. Main reason is that knowledge acquisition tasks and user oriented ergonomics rules compliance must be integrated in the overall engineering cycle, as we try to show it in the V-cycle diagram.

## 9. REFERENCES

[AGARD 1993 ] : "Knowledge Based System Approaches for Guidance and Control", AGARD Lecture series published in spring 1993.

[Air&Cosmos19934 ] : Air&Cosmos/ Aviation International N° 1480/81 11-24 july 1994.

[Amalberti et al 1990 ]: Amalberti René & Deblon François : "Cognitive modelling of fighter aircraft process control : a step towards an intelligent onboard assistance system", Int. J. of Man-machine studies.

[Aubry et al 1988 ] : Aubry P., Bracq D, Champigneux G., Gaudry Y., Havre A. : "Expert system model designed for airborne use on fighter aircraft", Human-Machine Interaction & A.I., Toulouse 1988.

[Brunessaux et al 1992 ], Jean Christophe Le Mentec and Stéphan Brunessaux , "Improving Reactivity in a Blackboard Architecture with Parallelism and Interruptions", Proceeding of the 10th European Conference on Artificial Intelligence, ECAI92, Vienna, Austria, August 3,7, 1992

[Champigneux et al 1989 ] : Champigneux G., Gaudry Y., Gilles A., Nory P., d'Orso M. : "Vers une intelligence artificielle embarquée pour l'assistance au pilote : le copilote électronique", 9_th International Conference on Expert Systems and AI, AVIGNON89, AI and Defense conf.,EC2, 1989.

[Desard et al 1991 ] : Desard F., Brunessaux S. , Bertuzzi S., Joubert T., Lementec J.C., Lalanda P., Charpillet F., Haton J.P., "Mécanismes temps réel dans un système multibases de connaissances", 11_th International Conference on Expert Systems and AI, AVIGNON91, AI and Defence conf,EC2, 1991.

[Dieng 1990 ] : Dieng Rose: "Méthodes et outils d'acquisition des connaissances", ERGO-IA congress ,Biarritz (Fr), 1990.

[Ermine 1991 ] : Ermine Jean-Louis : " Aspects méthodologiques et industriels des systèmes à base de connaissances", Journées état de l'art du GRECO-CNRS, 1992.

[Hickman et al 1989 ] : "Analysis for knowledge based systems. A practical guide to the KADS methodology", Ellis Horwood Ed., 1989.

[Joubert et al 1990 ] : Desard F., Collet M. ,

Bertuzzi S., Joubert T., Lementec J.C., Lalanda P., Charpillet F., Haton J.P., "Etude d'un système d'aide au pilote pour l'autoprotection d'un avion de combat", 10_th International Conference on Expert Systems and AI, AVIGNON90, AI and Defence conf,EC2, 1990 .

[Morillon et al 1988 ] : Morillon D., Conter T. & de Crémiers M. : "SEAN : an expert system for navigation assistance aboard fighter aircraft", 8_th International Conference on Expert Systems and AI, AVIGNON88, AI and Defence conf,EC2, 1988 .

[Rouse 1991 ] : Rouse William : "Design for success : A human centered approach to designing successful products and systems", Wiley series in systems engineering, 1991

[Sallé et al 1990 ] : Sallé S.E., Horak R., Berton D., Dagard L. "Realisation of a supervisor of an extended Kalman filter with application to satellite navigation" 11_th IFAC World Congress, Tallin (Est.) 1990

[Sallé et al 1993 ] : Sallé S.E., Champigneux G., Reichert O., Gaudry Y. "Development Methodology for Knowledge-Based Systems used in Guidance and Control : application to the COPILOTE ELECTRONIQUE project" , 13_i_th International Conference on Expert Systems and AI, AVIGNON93, AI and Defence conf,EC2, 1993.

[Smith et al 1988 ] Smith D.M., Broadwell M.M. "The pilot's associate : an overview", Avignon 1988

------------------------------------------------------------

# GROUP DISCUSSIONS

## 1. INTRODUCTION

Group Discussions were convened during the final two days of the meeting for the purposes of identifying, and of developing understanding on, new issues. Participants were divided into multi-disciplinary groups, each with a designated leader, and with a set of pre-selected issues to consider and to evaluate. This section summarises the reports of the group leaders to the final plenary session of the meeting.

## 2. PROCEDURE

The programme sought to make the workshop productive by developing and sharing understanding on new issues. Accordingly, the aim of the Group Discussions was to enable the emergent issues to be addressed in a systematic, and yet informal manner, with the widest participation from all attendees. In order to facilitate this process, all participants were invited to accept ownership of a specific key issue for the duration of the meeting, for discussion with others, and for reporting to the Group Discussions.

### 2.1 Issues Identification

Participants were directed to address areas of uncertainty, or problems requiring resolution. They were encouraged to present issues in the form of questions using the simple imperatives *Why? What? Which? How? Who? Where? and When?* The meeting call notice identified a series of issues relevant to Human-Electronic Crew teamwork. These issues, and others that emerged during the paper presentations, in particular those arising from the keynote, were collated and displayed for the participants consideration. Participants were encouraged to add further issues to this list as the meeting proceeded. A total of 43 issues was identified in this manner.

### 2.2 Initial List of Issues

A list of these 43 initial issues is shown below . This is not an exhaustive list of all the issues arising at the meeting. There were insufficient resources to capture all the issues raised by the papers, as they were presented. The list is merely intended to be indicative of the process followed by the meeting participants.

#### 2.2.1 *Call Notice Issues:*

(1). Do current development activities address the teaming issues?
(2). Are there some types or categories of decisions or actions that the Human-Electronic Team should never be trusted with?
(3). What oversight checks should be placed on the Team?
(4). How does the Team communicate with the higher authorities?
(5). Are there other issues besides teaming which are crucial to the operational application of the Electronic Crewmember concept?

#### 2.2.2 *Keynote Issues:*

(6). Who should be the team leader - the mission computer or the human?
(7). What types of teams should the system emulate?

(8). How do we ensure that the team samples we experiment on are representative?
(9). What human characteristics should we allow for in our team?
(10). How many humans should there be in the aircraft?
(11). How much should the team members trust each other?

#### 2.2.3 *Emergent Issues:*

(12). How do we distinguish reliably between true intentional and error behaviour?
(13). How do we get the pilot back in the loop?
(14). How do we inform the pilot about risk?
(15). What advice should not be given to the pilot in a given situation?
(16). Interoperability - How much?
(17). Co-ordination - How essential?
(18). Does machine advice promote habits or competence?
(19). How may the trust between the pilot and the decision aid be enhanced?
(20). How does one ensure that information of the highest priority is delivered to the pilot soon enough to make a difference to mission effectiveness and survivability?
(21). EC will require extensive data and software integration - what (im)possibilities for certification?
(22). For the near term, if we do have intelligent automation capable of successfully dealing with the caution, warning, advisory function, have we not captured the heart of the information processing architecture of the weapon system?
(23). Is it appropriate to reduce reliance on human decision making?
(24). How do we make uncertainty management effective?
(25). How do we elicit valued knowledge and eliminate reluctance to share knowledge?
(26). How do we validate knowledge about prospective systems, not yet in flight operation?
(27). Do we agree on the principles underlying case-based reasoning?
(28). How do we measure and predict fault size?
(29). How do we make human consent part of the decision criteria?
(30). What are the important differences in inferencing mechanisms for determining a decision?
(31). What can we do about people not having a common language for understanding events?
(32). Once decisions are arrived at, how should the decision be presented to the operator? As decision data with an associated probability or as a conclusion with action directions? Or, by the computer effecting (carrying out) the decision?
(33). Is "Trust" in the electronic crew enhanceable through training? Are some training techniques more effective in instilling trust than others?
(34). Human teams have a "power distance" (ranking). Who will be bold enough to shut down a triple redundant EC?
(35). Electronic member or 2-member crew, or electronic tools of a 1-member crew?
(36). Though getting the EC to "think like the pilot does" is one way to earn trust, this may lead to two heads making the same mistakes. Wouldn't 2 parallel but different (and co-

ordinated) problem solving approaches be more effective? If so, what are the parameters of co-ordination?

(37). Short of actually taking the pilot out of the cockpit, there will be a (hopefully increasing) number of tasks the EC can do *better* than the pilot. How can policy makers (pilots and above) make correct and optimal decisions about when to give a task to automation, when to the human, and when to mix? How can they implement/enforce those decisions?

(38). Do human crews have any relevance to human-electronic teams?

(39). How do we de-brief the EC? Will "his" explanations increase trust? How will he cope with the argument in the pub afterwards?

(40). Is trust between the EC and the operator being overemphasised as an issue?

(41). How do we ensure that the EC remains useful when the platform is used outside its designed role?

(42). When are we going to decide what we are trying to build and stop re-inventing the EC wheel?

## 2.3 Issue Assignment

Towards the end of the paper presentations, each participant was invited to select or accept ownership of an issue for the remainder of the meeting. The list of initial issues was offered as a guide, from which individual selections could be made. Alternatively, individuals were encouraged to take up a different issue of their personal choosing. This approach was aimed in part to accommodate individuals who wished to pursue a favourite or "pet" issue. The main purpose was for the owner to develop understanding of an emerging issue through informal discussion with others, particularly with participants who had indicated knowledge of the topic, or who held different views. Approximately half the attendees undertook ownership of an issue, in this way. Forms were

provided for recording information on the assigned issue. A worked example was demonstrated. This form requested information on the following:

(1) A statement of the issue.

(2) Implications for Human-Electronic Crew teamwork including AI technology, and the cockpit.

(3) Factors influencing the issue (AI & Cockpit).

(4) Other related issues (AI & Cockpit).

(5) Relevant knowledge, including lessons learnt, current practice, methods and techniques (AI & Cockpit).

(6) Potential directions, i.e. R&D requirements, alternatives, choices, priorities and cost/benefits (AI & Cockpit).

## 2.4 Group Assignment

After the paper presentations, the participants were divided into six multi-disciplinary groups, selected to enable a wide ranging discussion of the issues arising. The majority of the participants contributed to the group discussions. The aim of the discussion groups was to consider the assigned issues brought to the groups by the individuals, and any other issues, as the groups saw fit. In particular, the groups were directed to agree the priority of issues, and to provide justification for their decisions. A leader was assigned to each group, to manage the proceedings in accordance with the plan, and to record and to report on the group's deliberations. Forms were provided to structure the record of the decisions concerning issue priority and justification. These forms are illustrated in Figure 1 below. After group discussions, group leaders presented their conclusions during short 10 minute briefings in the plenary session of the meeting, with the aid of transparencies of the priority and justification forms.

| ISSUE | STATEMENT | PRIORITY L / M / H |
|---|---|---|
| 1 | | |
| 2 | | |
| 3 | | |
| 4 | | |
| 5 | | |
| 6 | | |
| 7 | | |
| 8 | | |
| 9 | | |
| 10 | | |

| JUSTIFICATION FOR HIGH PRIORITY ISSUE | |
|---|---|
| ISSUE STATEMENT : | |
| Implications of the issue for Human-Electronic Crew Teamwork | |
| Factors influencing the issue | |
| Other relevant issues | |
| Relevant knowledge,i.e. lessons learnt, current practice, methods and techniques | |
| Potential directions, I.e. requirements, alternatives, choices,priorities & cost/benefits | |

Figure 1   Forms used for reporting priority issues.

# 3. GROUP DISCUSSION REPORTS

## 3.1 Group 1

### 3.1.1 *Membership*
Tom Aldern (Issue: *How do we debrief the EC?*)
George Brander (Issue: *If we have not agreed models of human decision-making, how can we hope to build a PA which needs to know not only what the pilot needs but how he is thinking?*)
Charles Holly (Issue: *Where is the boundary between good human-computer interface design and technology?*)
Ian Ricketts (Issue: *Why is the technology not being shared with special needs research?*)
Tim Southam *Leader/Recorder* ( Issue: *To implement an EC, shouldn't we adopt a stepped approach, with aircrew involvement, to improve current and new cockpit modalities?*)

### 3.1.2 *Priority*
Four issue statements were generated by the group. These were then given a priority of low, medium or high. Three of the issue statements were given high priority, namely: To implement an EC, should we not adopt a stepped approach to improve existing / new cockpit modalities?; Can we justify *not* sharing technology? (N.B. PA = Pensioner's Associate!); Without models of human decision making can we build adaptive, intelligent systems? The group gave a low priority to the following issue statement: How do we debrief the EC to improve H-EC communication trust (trust is built up over years!) ?

### 3.1.3 *Justification*
a. Stepped Approach (*High Priority Issue*): It was agreed that implementing an EC should adopt a stepped approach to improve existing and new cockpit modalities. Improved data handling, improved modality, sensor fusion, high resolution displays, Head-Mounted Displays (HMD's), Head-Up Displays (HUD's), and Head-Down Displays (HDD's) were all discussed as having implications for H-EC teamwork. Factors which were identified by the group as influencing the issue included data requirements, system latencies, Man Machine Interface (MMI) input/output systems, and crew performance. It was agreed crew have the final say. Other issues highlighted as relevant to the present were technology, risk reduction, and simulation, particularly the need to replicate real world bad weather operations. On relevant knowledge, it was felt that progress should be made with what we know already. There was a need to define sensible specifications. Simulations should involve aircraft performing current tasks. In implementing an EC, increasing assistance as a stepped function, was seen as an important way forward incorporating new large HDD's and HMD's .

b. Sharing Technology (*High Priority Issue*): The issue of justifying *not* sharing technology, in particular with reference to special needs research, was addressed by the group. The implications for HE-C teamwork of sharing technology included multiplying your investment, access to practical solutions, access to experience, and access to tested novel approaches. Influencing the issue was risk i.e. it has worked before, and self interest (you will be old, and we will have quicker solutions). Other relevant issues included security (it has been handled!), the influence of teachers and graduates, and the need to visit the special needs group at Dundee. It was agreed that industry (e.g. AT&T, GEC, DEC) obviously use the relevant knowledge, but there was uncertainty about the use made within MoD, DoD, BAe etc. Potential directions included test projects, and a plea for the next project to involve Dundee (only 2.5 hours from London, 3.5 hours from Europe, 9 hours from America!).

c. Inadequate Models (*High Priority Issue*): The group considered that without agreed and understood models of human decision making, we cannot hope to build intelligent, adaptive systems. Humans can change strategies, but we do not know how. Therefore, how can make a Machine/Computer adapt to a human? Keeping knowledge simple results in keeping the Machine/Computer function fixed. Relevant factors included the lack of relevant cognitive models, and the continuing disagreement between classical decision-making theory and naturalistic approaches. A related issue was that we do not yet know whether the computer should think "like a human", or merely communicate in a "human-like" manner. Cognitive psychology and computer modelling were sources of relevant knowledge. Agreed and understood models of human decision making were seen to be needed for the implementation of an intelligent adaptive system. The group proposed further cognitive research as a way forward. They identified a need to field more prototypes, with further research thought to be needed on teams, perhaps using command cells and operations rooms (N=30) as analogues.

d. Debriefing (*Low Priority*). Consideration of debriefing the EC to achieve team co-ordination, and to build team trust, lead the group to questioning whether or not EC should be held to a higher performance standard than humans. Improvements in crew co-ordination between humans are achieved over years. It was questioned whether or not it was reasonable to expect EC to perform in an ideal fashion when humans would still make mistakes? Related to this issue were the problems of verification and validation of mission critical software. Also, for debriefing, it would be necessary to identify traceable actions from EC that can be recorded and reviewed post-mission. The Pilot's Associate programme was identified as a source of relevant knowledge, as well as experience gained with mission support tools. It was considered that we must provide a post-mission method to explain EC actions. Trust and confidence in human-human co-ordination improves over *many* missions, and Human-EC team co-ordination should be no different.

## 3.2 Group 2

### 3.2.1 *Membership*
Edmund Brugger
Jeremy Clare *Leader/Recorder*

Doug Hoecker
Chris Miller (Issue: *What authority structures / EC role in multi-user environment?*)
Gordon Semple: (Issue: *Electronic member of 2-member crew, or electronic tools of a 1-member crew?*)
Brian Sherwood-Jones: (Issue: *How do we provide guidelines on design and test procedures for the team as a unit?*)

### 3.2.2 *Priority*
The group identified the need to demonstrate the benefits of the Electronic Crew as a high priority issue. Other high priority issues were questions concerning the ability to measure and predict success of tools and EC systems, the required authority structure and EC role in the multi-use case, and how to provide design and test procedures guidance? Issues considered to have lower priority were as follows: Whether EC should replace or assist?; How safety is ensured?; How do we justify EC to society?

### 3.2.3 *Justification*
a. Demonstrable Benefits (*High Priority Issue*): It was agreed that there was a need to demonstrate as a community the benefits of EC. The implications of failing to demonstrate the benefits were that it would prevent the community ever building one. Factors influencing this issue were identified as being user requirements and technology capabilities. The group agreed that key problem areas for EC benefits were inadequacies in knowledge of the operational context and of human capabilities. Analysis of cost/benefits and options were potential ways forward.

b. Authority Structure(*High Priority Issue*): The implications of misunderstanding the issues of authority structures in a multi-user environment were that we would build the wrong system. For AI technology this was considered to have implications for task tracking, intent inferencing, planning, information management, task allocation and adaptive aiding. In the cockpit, the aim should be to provide each human with associate-like aiding, whilst facilitating co-ordination and minimising conflict. There was a need to understand authority in roles more widely. This included consideration of the "chain of command", the "definedness" of authority spheres, belief structures and practices in the domain, organisation, and individual crews, and trust between humans and machines. Social psychology and human communication, including Crew Resources Management (CRM), teamwork studies, and management practices, were identified as areas where relevant knowledge may be gained. Function/task analyses were also relevant, including tools and procedures such as CREWCUT, MIDAS, and the GEC Human Factors Task Data Base reported at the meeting. The conference paper on CASSY described a good working example. For the future, prototypes of multi-user systems should be developed in order to understand EC roles.

c. Measurement of Success (*High Priority Issue*): The question of how EC success was measured / predicted had implications for the design of decision aids, and for salesmanship. Factors influencing the issue were identified

as EC role, behaviour, authority structure and supporting technologies. Knowledge of system performance and CRM evaluations was relevant. Research was needed to evaluate various EC designs and role structures against multiple success measures, using novel evaluation approaches.

d. Design and Test Procedures. (*High Priority Issue*): Provision of guidance on design and test procedures for the team as a unit would mean that we would know how to have a design process that will end up with teamwork, and that we would know how to measure teamwork. Currently, we do not have methods for designing and testing HE-C cognitive systems. We need to develop and test design guidelines, trial procedures, and measuring instruments for HE-C teamwork as a whole. (N.B. This issue was submitted, but not discussed).

## 3.3 Group 3

### 3.3.1 *Membership*
Richard Bickerton (Issue: *How do we ensure the human-computer team functions when faced with unpredicted scenarios?*)
Peter Jorna *Leader/Recorder* (Issue: *Mode awareness or what will the thing be doing? Adaptive versus adaptable working environments?*)
Doug Hall
Tim Hughes
Taff Morgan ( Issue: *If your task analysis is of the wrong task, how can you use it to build up an EC to do the right task?*)
John Zenyuh (Issue: *What are the implications of moving EC technology towards a commercial environment?*)

### 3.3.2 *Priority*
Three issues of high priority and three of medium priority were identified by the group. Issues given high priority included: The study of operational persons in mission context; Design framework and the role of man; The need to cover the unexpected mission. Medium priority issues included: Moving EC technology into the commercial sector; EC and crew selection; Weapon deployment aid; Transfer to "full stereo".

### 3.3.1 *Justification*
a. Lack of Pilot-Vehicle Interface (PVI) Design Framework (*High Priority Issue*): It was considered that the principle requirement of crews is for 'information systems'. This did not automatically imply an EC. EC was only one way of presenting information to the crew, and not necessarily the most effective way, as shown by the video's on PA and CASSY. Research should be more focused on three aspects:

(1) The information needs of the crew
(2) The assignment strategy for functions and tasks
(3) The most effective human interface concept

These conclusions applied to any type of 'advanced automation' application e.g. air traffic management. There was insufficient emphasis at the meeting on providing an effective framework for the human interface design. Effort was needed in defining a basic design philosophy, with the

goal of preventing fragmented efforts within development programs that lacked funding for the crewstation. As a result of the lack of a framework, pilots and others have to cope with complex confusing designs and inconsistencies. These problems will become critical as demands increase arising from new missions. Advancement of the 'pilot vehicle interface' was needed to accommodate the opportunities provided by modern sensors and systems. Human factors had not kept pace because human factors knowledge was not integrated with knowledge of systems and operational requirements. Commercial organisation lacked the right combination of resources (versatile specialists and system oriented human factors group) to fill the gap. As a result, cockpit designs were driven by hardware and software and not by usability, and by the liveware. A project on 'Advanced PVI' was needed with the goal of providing a framework for design and evaluation, with a generic specification for a crew station to be customised by any manufacturer to their particular needs without leaving the basic validated concept.

b. In-Flight Studies (*High Priority Issue*): The group stressed a need for more detailed and structured 'in-flight' studies of pilot behaviour and task strategies in order to tailor the human interfaces. A design frame work requires studying basic concepts like 'adaptive' automation (system assumes tasks) versus 'adaptable' automation (pilot assigns tasks).

c. New missions (*High Priority Issue*): Future missions will be highly unpredictable. Thus, tailoring according to limited mission context will be out of date. The ability to adopt new mission capabilities should be a major design aim.

d. Weapon Deployment Aid (*Medium Priority Issue*): A first candidate for a pilot support system could be a function that provides feedback on the hit probability of specific target(s), depending on the characteristics of the weapon concerned. Not all weapons behave similarly, so there is no consistency for the pilot, and no predictability.

e. Commercial Technology Transfer (*Medium Priority Issue*) Moving EC technology into the commercial environment will require that there is appropriate Input/Output technology for EC in existing flight decks. For AI technology this will be influenced by market exposure, and by the risk perceived by both the public and the validation authorities. It will be influenced also by the opportunities for cockpit upgrades and for new systems. Other relevant issues include the extendibility of knowledge bases, and the construction of comprehensive knowledge bases. To achieve this will require knowledge of real-time AI techniques, and the ability to integrate with current D/C architectures. We will need to implement EC functionality with certified software languages and techniques. Also, we need to present legitimate cost/benefit arguments to the relevant decision-makers.



Figure 2. Implementation framework for user functional requirements.

### 3.4 Group 4

#### 3.4.1 *Membership*
Jim James
Richard Lynch (Issue: *What are the barriers to our problems with designing an effective pilot/EC interface, and how do we get around them?*)

Rick Magaldi
Tom Metzler *Leader/Recorder*
Michael Reinecke
Annette Simpson (Issue: *Is trust an important issue?*)

### 3.4.2 *Priority*

The need was recognised for a better understanding and knowledge of the barriers to, or problems with, designing an effective pilot / EC interface, along with solutions. However, the problem of identifying barriers to achieving effective PVI implementation was considered to raise a number of fundamental questions. In order of priority, a number of key issues were identified by the group. The highest priority question asked 'Where are we going?' Other questions, in reducing priority, were as follows: What are the stages?; What is the standard?; How much inter-operability? Consideration of the role of aiding in systems design, and of effective PVI implementation, followed from these issues.

### 3.4.3 *Justification*

a. Identification of Barriers *(High Priority Issue)* The need to identify the barriers to designing an effective pilot/EC interface, and how to get around them, was considered to have implications for the specification and prediction of EC capability, and the anticipated role of AI. In the cockpit, this issue has implications for the exploitation of avionics technology (flat panels, voice, pointing mechanisms). Influencing factors from AI technology included disagreements in philosophy concerning EC roles and functions, and reliance on metrics for determining requirements. In the cockpit, barriers included the use of limited symbol sets, graphics processor capability, transport delays, and the constraints of existing architectures. Other relevant issues were the general priority of AI development in Defence budgets, and negative press opinions of the PA programme. Related operational considerations included improvements in mission effectiveness and survivability, and improved flexibility, such as through mission replanning. Relevant knowledge for designing an effective interface should be derived from principles for intuitive design, from information requirements based on mission and segment analyses, and from the results of relevant studies, such as PCADS, ICAAS etc.

b. Where are we going, etc? *(High Priority Issues)* The justification for more fundamental issues was provided with reference to the diagram shown in Figure 2 above. It was considered that if the answer to the question "Where are we going?" was "to achieve implementation of the Pilot Vehicle Interface (with total AI end capability, and with resolution of the question of is the pilot or PA in charge)", then this needs to be defined now. This should be done by asking "What is the required design standard?". This will provide an aiming point. Then we need to ask "What is affordable for retrofitting (over 5/10/20 years), with reference to cost, risk and reliability?". Achieving interoperability and commonalty will require staged processes, including interfacing (between equipment particularly sensors, command elements, different systems and airframes), a common interface platform (translator / interpreter function), and a modular systems architecture. The role of aiding concepts in the design of systems needs to be defined to determine appropriate operator functions for successful PVI implementation.

### 3.5 Group 5

#### 3.5.1 *Membership*

Steve Boehmer *Leader/Recorder*
Robert Ellis (Issue: *How can we close the gap between the operator and system situation assessment?*)
Ian Burrett (Issue: *How do we select representative samples of humans for use in HE Crew research, and as the model on which to build systems?*)
Steve Selcon (Issue: *Why no progress in EC development in the last 4 years?*)
Mike Busbridge (Issue: *Just as trust is built up over time should the PA be taken a step at a time with perhaps only the system's health expert initially incorporated?*)

#### 3.5.2 *Priority*

Three high priority issues were discussed. The group considered as important the question of whether implementation should start slowly and incrementally. Also, concern was expressed about the apparent lack of progress in developing an EC in the last four years. A further issue of high priority concerned the question of selecting representative samples of humans for use in H-EC research and as a model on which to build a system. Other issues raised included: the gap between operator and system situation assessment; the need for EC requirements and specifications to be influenced by the user; and the importance and meaning of trust.

#### 3.5.3 *Justification*

a. Lack of Progress *(High Priority Issue)*. The apparent lack of progress in developing an EC in the last four years implied that there was still no coherent approach, philosophy, or design aims to apply to H-EC teamwork. Factors influencing the issue were identified in a lack of relevant basic psychology. The discussion prompted the group to reason there is a lack of continuity in research, and a lack of operator definition. The group felt that there was much reinventing of the wheel as far as relevant knowledge in progressing the development of an EC. If building the whole system was what was desired, then this should be the clear aim. But if this was not a practical option, then a less ambitious, and more pragmatic approach should be adopted.

b. Incremental Approach *(High Priority Issue)*. The need to start building EC slowly and to proceed incrementally was important in order to build up trust slowly. This approach was necessary because a full EC was not yet technically possible, and because the concept was a source of anxiety for pilots. Also, the high cost of financing the necessary work made full EC unaffordable. The first step of this incremental approach probably should be to build a systems health expert. Cartoons are provided below, in Figure 3, to illustrate this argument (courtesy of Mike Busbridge).

c. Representative Samples *(High Priority Issue)*. Selecting the right crews for teamwork research was important because of the danger of optimising for the wrong model. A wide range of performance might be achieved with such close

interaction. The teamwork model should exploit human strengths, and accommodate weaknesses, and perform in a predictable manner. Trained aircrew subjects, with appropriate experience in certification test and evaluation, have limited availability, but they can provide important "non-parochial" feedback. Data was needed on the range of human performance achieved, coupled with objective and subjective workload measures. This would enable banding of the range of human characteristics that work effectively as an H-EC, which could be fed into aircrew selection and training.

d. Situation Assessment (*Medium Priority Issue*). The existence of a gap between operator and system situational awareness meant that if the situation representation was poor in the system, then incorrect decisions could be made, or correct decisions arrived at for the wrong reasons. Mismatch between operator and system assessment would reduce trust in system actions. Relevant AI factors included knowledge acquisition, informal interfaces, estimating operator intent, and knowledge representation. Input/output technology for airborne environments was considered to influence situation assessment in the cockpit. Cognitive interviewing techniques (e.g. critical incident methods) provided useful data. Prototypes could be used as platforms for knowledge acquisition. Focus was needed on operator-system integration. Integration was necessary for both situation assessment and planning activities. These should be thought of as shared rather than divided tasks. The dual aim should be to reduce workload and improve effectiveness. Input/output technology could be as influential as automation technology, in determining the outcome.



**Figure 3. Cartoons illustrating the justification for an incremental approach.**

### 3.6 Group 6

#### 3.6.1 *Membership*
Peter Crosby (Issue: *Without an achievable goal there is a danger of going around in circles. The line between limited information displayed to the pilot and a fully automated RPV is an experimental curve in terms of time and content. Where on the curve do we want to aim for?*)
Hansjorg Hertweck
Peter Raeth *Leader/Recorder* (Issue: *How does one ensure that information of the highest priority is delivered to the pilot soon enough to make a difference to mission effectiveness and survivability?*)
Jerry Seeman (Issue: *Should EC provide options or directions?*)

Iain Macleod (Issue: *Does machine advice promote habits or competence?*)

#### 3.6.2 *Priority*
The individual issues were all considered important, but through discussion a different set of issues emerged, with some consensus on priorities. Three high priority, three medium priority and two low priority issues were identified by the group. High priority issues included concerns about the lack of rigorous team definitions? The group thought it important to consider objective design and development criteria, and how an EC could be flight and combat certified? Questions about EC training, ensuring that EC works fast enough, and EC debriefing, were considered to be of medium importance. Issues given a lower priority included the level

179

of detailed information that a human uses effectively, and how to prevent the human from getting out of the loop?

### 3.6.3 Justification

a. HE-C Definitions, Criteria and Certification (High Priority Issues). Resolution of the high priority issues of providing rigorous H-EC team definitions, objective H-EC design and development criteria, and valid EC certification methods, will have implications for the structure and direction of research. It will exert a major influence on systems integration, and provide criteria and metrics for evaluation. Crew acceptance, funding and timing, and both hardware and software technology will influence H-EC implementation. Greater integration of the research community was required than at present, and appropriate evaluation methods needed to be made available. Areas of relevant knowledge include decision aid technology, human factors and psychology, validation and verification, and knowledge acquisition. Future directions will require the use of full parallel, concurrent and co-processing. Verification and validation methods will need to be extended. Dynamic function allocation will need to be provided to make sure that the crew stays involved. Event logging and analysis will be needed for EC systems. Specification methods will need to be developed for EC systems.

b. Information Delivery (High Priority Issue) It was considered important to ensure that information of the highest priority will be delivered to the pilot soon enough to make a difference to mission effectiveness and survivability. If such information could not be implemented so that it was fast enough to make a difference, AI technology would be abandoned. It was necessary to overcome the 'toy problems only' reputation. AI technology needed to be linked to the needs of operational crews. Modern cockpit tasks were complex with rapid response rate requirements. The result was task overload for reduced crews. AI technology could solve this problem but only if it was fast enough to make a difference. The task overload barrier was waiting to be breached. AI technology had the potential to be successful but it needed to be linked to the needs of operational aircrew. Relevant AI factors included the following: processor speed; data bus speed; software and hard-ware architectures; scaleable implementations; parallel, concurrent and co-processing; knowledge and example acquisition; validation and verification. Cockpit factors influencing this issue were as follows: pilot control of EC; provision of the appropriate level and volume of information; sufficiently high screen resolution and refresh rate; pilot trust in EC; failure rate and accuracy of EC. Other relevant issues included co-ordination of on-board with off-board information, and in the cockpit, the sunlight readability of colours and screens, and laser protection. It was considered essential to employ appropriate computer programming languages from the start, and not use methods which require "black magic" for successful implementation. Current practice should be employed in the areas of neural nets; network, rule, matrix-based expert systems; fuzzy logic. Trust will not be easily recovered in high complexity, high risk environments when EC was known to fail or to be inaccurate. For the future, work should seek to merge neural nets, expert systems and fuzzy logic since the math was similar and appeared to be easily implemented via optics. As regards the cockpit, work should pursue objective studies of trust. Dynamic function allocation should be applied to obtain the appropriate human workload.

c. Options or Directions (Medium Priority Issue) The question of whether EC should provide options or directions was considered to have implications for the generation of probabilistic calculations, and for the provision of directional displays in the cockpit. The issue was thought to be influenced by the data sources available, data validity, and the availability of crucial data. In the cockpit, influencing factors would be the acceptance of risk by the operator, and the changing risk over the course of the mission. Other relevant issues included the debate about data or information provision, and the necessity of providing explanations to the operator. Relevant knowledge identified included soft computing methods, and hybrid "soft" and binary systems. In future, work should consider how to combine computational systems for useful outputs, and how to address the development of effective display formats.

c. Complacency (Medium Priority Issue) The question of whether machine advice promotes habits or competence was considered to be of low priority for systems that are properly designed, with the pilot in control of EC, and with EC in an assistant function. But in poorly designed systems, with the pilot in a supervisory role, machine advice could lead to complacency, and to inflexible non-adaptive behaviour. Since it seemed generally accepted that "pilot in control" was the intended design aim, then complacency ought not to be a major problem. A schematic diagram was used to illustrate the intended "pilot in control", as shown below in Figure 4



**Figure 4. "Pilot in Control" Model**

180

# 4. CONCLUSIONS

If there is a simple answer to the question "Can we trust the team?" then it may be "Yes, providing we sort out these issues". A list of the issues considered by the Discussion Groups is shown below in Table 1. The issues are arranged in order of the priorities assigned by the Discussion Groups, beginning with the high priority issues. The number of the Discussion Group addressing the issue is indicated to facilitate cross-reference to the preceding text.

Some general points arise from this analysis. Understandably, the process was biased towards identifying high rather than low priority issues. Groups 1, 2, 3 and 5 mostly evaluated and prioritised the issues brought along by the individual group members. Groups 4 and 6 sought more to present an integrated overview, moving on from the individual issues. The framework provided by Group 4 seems particularly helpful. Group 6 identified a key working assumption, governing the outcome of many issues, i.e. that the pilot needs to be, and will be, in control. Some of the issues given high priority reflected particular interests of influential individuals, and these may have reduced consensual validity as a result. But most were relevant to the proceedings, and directly associated with the aims of the meeting. Assessments of issues with a either high cockpit/crewstation AI content, and/or important cockpit/crewstation AI implications, probably have the greatest consensual validity.

Generally, the groups seemed concerned to highlight strategic issues and governing processes, rather than specific, detailed technical issues. Most groups emphasised the issues with a major bearing on the implementation of H-EC concepts and requirements. Examples include the definition and setting of objectives; establishment of design frameworks, guidelines, standards and criteria; incremental development strategies; and measures of effectiveness and operational validity. This may be because the multi-disciplinary nature of the groups resisted conscensus on priorities between technical areas. Also, since the individual papers dealt with technical issues in depth, it may have been appropriate to highlight the "big issues" at this point in the proceedings.

Some "big" technical issues were highlighted, such as the inadequacies of human decision-making models, the need to establish appropriate role and authority structures, principles for selecting experimental samples, valid certification procedures, and achieving effective functioning for the unpredicted mission. None of the above could be regarded by all as an entirely new issue, in the sense of being novel, and unfamiliar. They could be more appropriately characterised as the burning issues of the day. Entirely novel issues are perhaps to be found embedded in the texts of the technical papers, and as yet largely unrecognised, except presumably by the authors.

During the wash-up session, the UK user representatives stated that a more effective combination would be achieved by retaining, not replacing, the inherent aircrew capabilities, and by combining with the AI technology to produce more effective systems. It was also suggested that better communication was required between the research community and the user community to increase awareness of the potential of AI technology. It was then pointed out that, in the UK at least, the surface, sub-surface, and airborne Navy were already pursuing ambitious programmes in exploiting AI technology, and so this awareness had already been achieved by one branch of the UK armed forces.

Understanding trust, perhaps not surprisingly, since it received such close attention in many of the technical papers, was not such a burning issue by the end of the meeting. Could this be a measure of the meeting's success?

| ISSUE | STATEMENT | GROUP | PRIORITY |
|-------|-----------|-------|----------|
| 1 | To implement an EC, shouldn't we adopt a stepped approach to improve existing and new cockpit modalities? | 1 | High |
| 2 | Can we justify not sharing technology? | 1 | High |
| 3 | Without models of human decision making, can we build adaptive intelligent systems? | 1 | High |
| 4 | What authority structures and EC roles should there be in a multi-user environment? | 2 | High |
| 5 | How do we measure and predict success? | 2 | High |
| 6 | How do we, as a community, demonstrate the benefits of EC? | 2 | High |
| 7 | How do we provide guidelines on design and test procedures for the team as a unit? | 2 | High |
| 8 | How do we provide a design framework to include the role of the operator? | 3 | High |
| 9 | How do we ensure that the H-EC team functions when faced with the unpredicted mission? | 3 | High |

| 10 | How do we that the design is based on the right operational tasks and context? | 3 | High |
|---|---|---|---|
| 11 | What are the barriers to our problems with designing an effective pilot/EC interface, and how do we get around them? | 4 | High |
| 12 | Where are we going? | 4 | High |
| 13 | What are the stages? | 4 | High |
| 14 | What is the standard? | 4 | High |
| 15 | How much interoperability? | 4 | High |
| 16 | How do we employ aiding in design? | 4 | High |
| 17 | How do we achieve effective PVI implementation? | 4 | High |
| 18 | Should we start slowly and incrementally? | 5 | High |
| 19 | Why has there been no development in EC in the last 4 years? | 5 | High |
| 20 | How do we select representative samples of humans for use in H-E crew research and as a model on which to build systems? | 5 | High |
| 21 | What are the teams and their rigorous definitions? | 6 | High |
| 22 | What are the objective design and development criteria? | 6 | High |
| 23 | How would an EC be flight and combat certified? | 6 | High |
| 24 | How does one ensure that information of the highest priority is delivered to the pilot soon enough to make a difference to mission effectiveness and survivability? | 6 | High |
| 25 | Should EC replace or assist? | 2 | Medium |
| 26 | How do we ensure safety? | 2 | Medium |
| 27 | How do we justify EC to society? | 2 | Medium |
| 28 | Why not first build a weapon deployment aid? | 2 | Medium |
| 29 | Why not transfer to "full stereo"? | 2 | Medium |
| 30 | How do we do EC and crew selection? | 3 | Medium |
| 31 | How do we close the gap between operator and system situation assessment? | 5 | Medium |
| 32 | How do we make EC requirements specifications more influenced by the "user". | 5 | Medium |
| 33 | Is trust an important issue and what do we mean by trust? | 5 | Medium |
| 34 | Should EC provide options or directions? | 6 | Medium |
| 35 | Does machine advice promote habits or competence? | 6 | Medium |
| 36 | How is an EC to be trained? | 6 | Medium |
| 37 | How do we ensure that an EC works fast enough | 6 | Medium |
| 38 | How does an EC get debriefed post-mission? | 6 | Medium |
| 39 | How do we debrief EC to improve Human-EC communication and trust? | 1 | Low |
| 40 | What level of detail can a human effectively use? | 6 | Low |
| 41 | How do we keep the human from getting out of the loop? | 6 | Low |

**Table 1. Final List of Issues.**

# ATTENDEES

Tom Aldern
GreyStone Technology, Inc.
15010 Avenue of Science Street, Suite 200
San Diego
California 92128
USA.
Tel: (C) 619-675-7800
Fax: (C) 619-675-7808


Dr Deborah Bekerian
MRC Applied Psychology Unit
15 Chaucer Road
Cambridge CB2 2EF
UK
Tel: (C) (0)223-355294
Fax: (C) (0)223-359062


Lt. Richard Bickerton
Man Machine Integration Department,
R184 Bdg
DRA Farnborough
Hampshire, GU14 6TD
UK.
Tel: (C) (0)252-392486
Fax: (C) (0)252-376392


Lt. Col. Steven Boehmer
US Air Force
Joint Cockpit Office (WL / XPK)
2210 Eight Street, Suite 1
Wright Patterson AFB
Ohio, OH 45433,
USA
Tel: (C) 513-255-6644
Fax: (C) 513-476-4547


George Brander
CSSS/2 Block 2
DRA Portsdown
Portsmouth
Hampshire, PO6 4AA
UK
Tel: (C) (0)705-333357
Fax:(C) (0)705-333543


Edmund Brugger
DASA-LM
Dornier
Luftfahrt GmbH
Postfach 1303
D-7990 Friedrichshafen 1
Germany
Tel: (C)-49-7545-84990
Fax: (C)-49-7545-84927

Sqn Ldr Ian Burrett
OR52a (Air),
Room 2249,
MOD Main Building
Whitehall
London SW1A 4AH
Tel: (C) (0)71-218-7776
Fax: (C) (0)71-218-7474


Michael Busbridge
Venom Management Team
Displays Group
GEC Marconi Avionics
Airport Works
Rochester,
Kent ME1 2XX
UK
Tel: (C) (0)634-844400
Fax: (C) (0)634-827332


Jeremy Clare
Cambridge Consultants Ltd
Science Park
Milton Road
Cambridge CB4 4DW
UK
Tel: (C) (0)223-420024
Fax: (C) (0)223- 423373


Anthony Clark
GEC Marconi Research Ltd
West Hanningfield Road
Great Baddow, Chelmsford
Essex CM2 8HN
UK
Tel: (C) (0)245-473331 Ext 3244
Fax: (C) (0)245-475244


Capt Peter Crosby
MOD (Army)
Development and Trials Sqn
Army Air Corps
DAAVN  Middle Wallop
Stockbridge
Hampshire SO20 8DY
UK
Tel: (C) (0)980-67-4284
Fax:(C) (0)980-67-4105


Jo Davies
British Aerospace (Military Aircraft) Ltd.
York House,
Farnborough Aerospace Centre
Farnborough
Hampshire GU14 6YU
UK
Tel: (C) (0)252-382282
Fax: (C) (0)252-382060

Dr John Dennett
MRC Applied Psychology Unit
15 Chaucer Road
Cambridge CB2 2EF
UK
Tel: (C) (0)223-355294
Fax: (C) (0)223-359062

Robert Ellis
Man Machine Integration Department,
R184 Building,
DRA Farnborough,
Hampshire GU14 6TD
UK
Tel: (C) (0)252-392486
Fax: (C) (0)252 392770

Marc Gerlach
Universitat Der Bundeswehr Muhchen
Institut fur Systemdynamik und Flugmechanik
Werne-Heisenberg-Weg 39
D-85577 Neubiberg
Germany
Tel: (C) 49-89-6004-2139
Fax: (C) 49-89-6004-2082

Gretchen Greatorex
GEC- Marconi Research Centre,
West Hanningfield Road,
Great Baddow,
Chelmsford,
Essex. CM2 8HN
UK
Tel: 0245-473331 Ext 3247
Fax: 0245-475244

Douglas Hall
Delex Systems Inc.
5100 Springfield Rd, Suite 200
Dayton
Ohio, OH 45431
USA
Tel: (C) 513-259-0231
Fax: (C) 513-259-0924

Dr Peter Hancock
141 Mariucci Arena-Operations,
1901 4th Avenue, SE,
University of Minnesota,
Minneapolis, MN 55455
USA
Tel: (C) 612-626-7521
Fax: (C) 612-626-7494

Hansjorg Hertweck
DASA - LM
Dornier
Luftfahrt GmbH
D-88039 Friedrichshafen 1
Germany
Telephone: (C) 49 7545-82463
Fax: (C) 49 7545-84927

Doug Hoecker
Human Sciences Group 401-2B4
Westinghouse Science & Technology Centre
1310 Beulah Road
Pittsburgh
Pensylvania PA 15235
USA
Tel: (C) 412-236-2750
Fax: (C) 412-256-1348

Charles Holly
Bell Helicopter-Textron Inc.
PO 842
Fort Worth
Texas 76101
Tel: (C) 817-280-4086
Fax: (C) 817-280-3138

Howard Howells
Man Machine Integration Department,
R184 Building
DRA Farnborough
Hampshire, GU14 6TD
UK
Tel: (C) (0)252-392481
Fax: (C) (0)252-376392

Tim Hughes
Westland Helicopters LTD,
Lysander Road,
Yeovil,
Somerset. BA20 2YB
UK
Tel: (C) (0)935 702287
Fax: (C) (0)935-703752

Jim James
AHUFT Block 6,
Army Air Corps Centre,
Middle Wallop,
Stockbridge,
Hants SO20 8DY
UK
Tel: (C) (0)980-674118
Fax: (C) (0)980 674110

Peter Jorna
Netherlands Aerospace Laboratory NLR
Anthony Fokkerweg 2
1059 CM Amsterdam
The Netherlands
Tel: (C)-31-20-5113638
Fax: (C)-31-20-5113210

Thierry Joubert
MATRA-Defense
DAST/SIA
37 Avenue Louis Breguet
78146 Velizy Villacoublay
Cedex
France
Tel: 34884612
Fax: 34884455

Angela Lucas
Logica UK Ltd
Betjeman House
104 Hills Rd
Cambridge CB2 1LQ
UK
Tel: (C) (0)223-66343
Fax: (C) (0)223 322315


Richard Lynch
McDonnell Douglas Aerospace,
P.O. Box 516
M/C 0642203,
St. Louis,
Missouri MO 63166-0516
USA
Tel: (C) (314) 233-4934
Fax: (C) (314) 233-4433


Rick Magaldi
Artificial Intelligence Manager
British Airways Information Management
Viscount House (E90)
PO Box 10
Heathrow Airport
Hounslow
Middlesex TW6 2JA
UK
Tel: (C) (0)81- 562-6933
Fax: (C) (0)81-562-4475


Tom Metzler
Army Aviation Troop Command,
ATCOM AMSAT-R-ESC,
4300 Goodfellow Building,
St. Louis,
Missouri MO 63120-1798
USA.
Tel: (C) 314-263-1130
Fax: (C) 314-263-1622


Iain MacLeod
Aerosystems International
West Hendford
Yeovil
Somerset BA20 2AL
UK
Tel: (C) (0)935-443014
Fax: (C) (0)935-34145


Christopher Miller
Honeywell Technology Center,
3660 Technology DR,
MN65-2600
Minneapolis, MN 55418
USA
Tel: (C) (612) 951-7484
Fax: (C) (612) 951-7348

Gp Capt G A Miller
DDOR5(Air)
Room 2231
MOD Main Building,
Whitehall
London SW1A 4AH
UK
Tel: (C) (0)71-21-81540
Fax: (C) (0)71-21-87474


Sqn Ldr Charles Morgan
SA1
Armament and Mission Systems Division
A&AEE Boscombe Down
Salisbury
Wiltshire SP4 OJF
UK
Tel: (C) (0)980-66327
Fax: (C) (0)980-663163


Dr Chris Ovenden
Smiths Industries Aerospace and Defence Systems
Bishops Cleeve
Cheltenham
Gloucestershire GL52 4SF
Tel: (C) (0)242-673333
Fax: (C) (0)242-661795


Major Peter Raeth
USAF Wright Laboratory
WL/FIPA,
2210 Eighth Street Suite 1,
Wright Patterson AFB
Ohio OH 45433-7511
USA
Tel: (C) 513-255-8252
FAX: (C) 513-476-4547


Michael Reinecke
Freiwilligenannahmestelle Ost
Regattastr 12,
12492 Berlin,
Germany.
Tel: (C) 49 30-63 047 400
Fax: (C) 49 30-67 64 216


Dr Ian Ricketts
Microcentre
University of Dundee
Dundee, DD1 4HN
UK
Tel: (C) 01382 344153
Fax: (C) 01382 23435


Dr John Reising
USAF Wright Laboratory
WL/FIPA,
2210 Eighth Street Suite 1,
Wright Patterson AFB
Ohio OH 45433-7511
USA
Tel: (C) 513-255-8252
Fax: (C) 513-476-4547

185

Peter Robertson
Loral-ASIC
Po Box 41
North Harbour
Portsmouth
Hampshire, PO6 3AU.
UK
Tel: (C) 0705-564168
Fax: (C) 0705-383546

Dr Graham Rood
Mission Management Department
R177 Building
DRA Farnborough
Hampshire GU14 6TD
UK
Tel: (C) (0)252-393076
Fax: (C) (0)252-376392

Pierre Sassus
SEXTANT Avionique
Aerodrome de Villacoublay - BP59
F78141 Velizy-Villacoublay
Cedex
France
Tel: (C)-33-(1)-46 29 71 87
Fax: (C)-33-(1)-46 29 77 66

Jerry Seeman
McDonnell Douglas Helicopter Systems
5001 E. McDowell Road
Meza
Arizona AZ 85207
USA
Tel: (C) 602-891-6707
Fax: (C) 602-891-6003

Dr Stephen Selcon
Man Macine Integration Department
R177 Building
DRA Farnborough
Hampshire GU14 6TD
UK
Tel: (C) (0)252 392155
Fax: (C) (0)252 376392

Gordon Semple
British Aerospace Defence
Military Aircraft Division
W392D
Warton Aerodrome
Preston,
Lancashire PR4 IAX
UK
Tel: (C) (0)772 633333
Fax: (C) (0)772-855065

Brian Sherwood-Jones
BAe Semma
1 Atlantic Quay
Broomielaw
Glasgow G2 8JE
Tel: (C) (0)41-204-2737
Fax: (C) (0)41-221-6435

Dr Annette Simpson
CSSS/2 Block 2
DRA Portsdown
Portsdown
Hampshire PO6 4AA
UK
Tel: (C) (0)705-332418
Fax: (C) (0)705-333543

Tim Southam
Man Machine Integration Dept.
Room 134
R177 Building
DRA Farnborough
Hampshire  GU14 6TD
Tel: (C)-(0)252-392897
Fax: (C) (0)252-376392

Alison Starr
Smiths Industries
Aerospace and Defence Systems
Bishops Cleeve
Cheltenham
Gloucestershire GL52 4SF
Tel: (C) (0)242-673333
Fax: (C) (0)242-661795

Robert Taylor
DRA Centre for Human Sciences
CHS 3
F131, Rm 43
Farnborough
Hampshire GU14 6SZ
UK
Tel: (C) (0)252 394120
Fax: (C) (0)252 392984

George Ward
ESE Associates
15 Jesse Close
Yateley
Camberley
Surrey GU1 AH
UK
Tel: (C) (0)252-876783
Fax: (C) (0)252-890259

Elizabeth Wheatley
DRA Centre for Human Sciences
CHS 3
F131, Rm 43
Farnborough
Hampshire GU14 6SZ
UK
Tel: (C) (0)252 394364
Fax: (C) (0)252 392984

John Zenyuh
Search Technology
4898 South Old Peachtree Road
Atlanta
Georgia GA 30071
Tel: (C) 404-441-1457
Fax: (C) 404-263-0802