Comparing Design Options for Allocating Communication Media in Cooperative Safety-Critical Contexts: A Method and a Case Study

ROBERT FIELDS
Middlesex University
FABIO PATERNÒ and CARMEN SANTORO
CNUCE-C.N.R.
and
SOPHIE TAHMASSEBI
C.E.N.A.

In this article we present a method for evaluating and comparing design options for allocating communication media. The method pays particular attention to how such options support cooperation in an interactive safety-critical system. The comparison is performed using three sets of criteria based on task performance, analysis of user deviations and consequent hazards, and coordination. The explicit emphasis on hazards and communication issues, using actual tasks to guide the evaluation, ensures that designers' attention is focused on the interactions where problems are likely to occur. We describe an application of the method to the design of access to new communication technology in an air traffic control environment.

Categories and Subject Descriptors: H.5.2 [**Information Interfaces and Presentation**]: User Interfaces—*Evaluation/methodology*; *Input devices and strategies*; *Interaction styles*; H.5.1 [**Information Interfaces and Presentation**]: Multimedia Information Systems—*Audio input/output*; D.2.2 [**Software Engineering**]: Design Tools and Techniques—*User interfaces*

General Terms: Design, Human Factors, Reliability

Additional Key Words and Phrases: Air traffic control, tasks, usability and safety

We gratefully acknowledge support from the MEFISTO European LTR Project (http://giove.cnuce.cnr.it/mefisto.html).

Authors' addresses: R. Fields, School of Computing Science, Middlesex University, Bounds Green Road, London, N11 2NQ, UK; email: B.Fields@mdx.ac.uk; F. Paternò and C. Santoro, CNUCE-C.N.R., Via V. Alfieri 1, 56010 Ghezzano, Pisa, Italy; email: F.Paterno@cnuce.cnr.it; C.Santoro@cnuce.cnr.it; S. Tahmassebi, C.E.N.A., Avenue Edouard Belin, 7, Toulouse, 31055, France; email: Tahmasse@cena.fr.

1. INTRODUCTION

1.1 Safety and Usability

In many safety-related systems, computer-based equipment is being used to augment a number of communication functions. This has the potential to transform the way people work by providing new facilities, and making existing ones available in new and more flexible ways. Media allocation design issues are becoming increasingly critical as technology becomes more diverse and more pervasive. *Media allocation* refers to decisions about the access that actors have to different communication media in a system. This is of particular significance in safety-critical systems where many studies have shown that accidents often are caused by a human error whose likelihood may be increased by poor design [Reason 1990]. For instance, Hollnagel [1993] surveys literature citing human error as a causal factor in as many as 80% of safety-related incidents across a range of high-technology industrial sectors.

In such "high-consequence" systems two contradictory tendencies seem to coexist. On the one hand, it is widely believed that the deployment of advanced technology can lead, by automating inefficient and error-prone tasks, to improvements in both performance and safety. On the other hand are the concerns that the introduction of new technology may have negative and unanticipated consequences, typically by adding complexity to systems in which the ability of humans to act effectively is of extreme importance.

In this regard we believe, that, instead of relying solely upon empirical testing in later phases of the development lifecycle, it is vital to perform evaluations of different user interface proposals and task and media allocation choices early in the development process.

1.2 Evaluation of Interactive Safety-Critical Systems

Despite the multitude of existing evaluation techniques in HCI, we note that few approaches provide an integrated and comprehensive analysis of the factors that play an important role in an interactive safety-critical application. GOMS approaches [John and Kieras 1996; Gray et al. 1993] mainly focus on performance efficiency and error-free behavior, but have paid little attention to other aspects of human-machine interaction. Inspection-based methods such as heuristic evaluation [Nielsen 1993] and cognitive walkthrough [Wharton et al. 1994] have been found useful for detecting basic usability problems without providing any particular support for those that can have a strong impact on the safety of the system. For example, both cognitive walkthrough and our method imply the analysis of a sequence of actual tasks in order to identify possible user problems. However, while cognitive walkthrough mainly aims at identifying whether or not the user will be able to perform the right interaction, our method is oriented to understand the effects that can be triggered when users perform the wrong interaction or perform the right interaction at the wrong time. This is important for suggesting an improved design that can either prevent user deviations or limit their possible impact on the safety of the system.

In addition, in the safety engineering area many techniques for exposing safety-related problems have been developed, including HAZOP [UK Ministry Of Defence 1996] and human reliability analysis [Kirwan 1994], but there is the problem to find an effective way to apply such methods so as to improve the design of the user interface. For example, HAZOP techniques have mainly been used to identify hazards arising from problems in the internal functioning of a safety-critical system. Human reliability analysis stems from the need to quantify the effects of human error on the risks associated with a system, and typically involves estimating the probability of the occurrence of errors. In practice, the assessment of human reliability very soon runs into the problem of acquiring the data needed for reliable quantification. Practical experience shows that different methods in this area often yield different numerical results, and even the same method may give different results when used by different analysts [Hollnagel 1993].

MECHA (Method for Evaluation of Cooperation, Hazards, and Allocation) seeks to remedy some of these problems by incorporating three key elements: it requires the development of task modeling and the identification of important scenarios of use; it analyzes possible user "deviations"; and it considers the cooperation that occurs between different individuals in the system. This provides designers with the necessary elements for an integrated qualitative analysis of usability and safety aspects.

Safety, as many authors have observed, is not a property of individual tasks or actions, but of the interrelationships and interconnections between parts of a system. Perrow [1984], for instance, identifies the complexity of a system and the coupling between its parts as significant factors in the genesis of accidents. What is often referred to as "human error" cannot be seen simply as a result of failures in human information processing; technology and technical change in a work system can create contexts that shape the way actions—erroneous or otherwise—take place [Woods et al. 1994]. MECHA supports the investigation of competing design proposals in terms of how they might encourage or discourage various types of failures, how design and allocation decisions may tend to mitigate the effects of failure, and the ways that technology contributes to the detection and repair of failures. In the context of this work, we are particularly interested in safety issues in the communication between users of *a collaborative system*. Air traffic control work involves the coordination of activities and reconciliation of interests of pilots, the controllers working within a sector, and the controllers of other sectors, as well as many other agents and agencies. Symon et al. [1996] argue that in complex systems a number of conflicts may exist in such work (for instance, between formal constructs and the work goals, or between the goals themselves), and that the term "collaboration" may be misleading.

It must be emphasized that technology changes inevitably lead to changes in the tasks of individual actors in the system, a fact reflected in

Carroll and Rosson's [1992] discussion of the "task-artifact" cycle. In particular, any device may provide support, automate tasks, or solve problems for some aspects of a person's activity. At the same time, the introduction of technology may have the effect of transforming tasks, creating new demands, and placing additional requirements on the humans in the system.

Finally, a need for a structured analysis has arisen from the consideration that ethnographic approaches, by themselves, are not sufficient to provide the kind of information required by system designers and developers [Bentley et al. 1992]. Ethnographic approaches tend to provide many details without sufficient indications about priorities among them. Such a sea of details can complicate the work of designers, who often have to make decisions about which details are most relevant. MECHA is a method for structured analysis that can be applied to a variety of systems, not only to those that are safety critical. However, some aspects of the method, in particular the analysis of deviations and hazards, are aimed mainly at systems where safety is a prime concern.

1.3 The Air Traffic Control Domain

Air traffic control (ATC) systems are highly interconnected cooperative systems that face constant pressure to innovate. Huge increases in the volume of air traffic have meant that the existing systems are beginning to find it difficult to cope. The results are not only delays for the traveling public, but also concerns about the risks of near misses and other incidents. A typical response is to meet the increased traffic levels and consequent increases in controller workload with proposals for increased technological sophistication. Previous attempts have been made to develop alternative user interfaces for controllers [Chatty and Lecoanet 1996] or to augment the existing environment with novel interaction techniques [Mackay et al. 1998]. However, despite much research effort, this work has remained at a prototype stage and has not been deployed in the field.

One of the problems currently facing the ATC system arises because most communication between controllers and pilots is carried out using VHF radio, a medium of limited bandwidth that is fast becoming a bottleneck. This limitation, together with the known failure modes of voice communications (misheard communications being one of the most commonly reported ATC problems), can have serious safety implications. One solution that has been proposed is the introduction of *data link*, a technology allowing asynchronous exchanges of digital data coded according to a predefined syntax. This technology seems to overcome some of the main limitations suffered by the traditional system, but its implications for the work of controllers and the safety of the overall system are not fully understood.

1.4 Structure of the Article

The purpose of this article is twofold. First, we describe how the MECHA method aims to help designers to understand and analyze the impact of new communication technology in a safety context, and the different ways of assimilating such technology into the existing working practices. Second, we illustrate the application of MECHA to a real case study: the introduction of data link technology into en-route air traffic control. We discuss only a part of the domain, though the method can be applied much more widely. We consider a set of tasks and a scenario, and we use them as a means to compare and evaluate competing design options.

In Section 2 we introduce the method, and in Section 3 outline the case study in general terms. Section 4 moves on to describe the application of the method to analyzing possible problems, activities, and issues in the current system. Next, we propose a different way to allocate the data link and flight information to the controllers in a sector, and thereby allocate the main tasks between them (Section 5). Then, we summarize our results and discuss how the integrated analysis of task, technology, and scenarios performed through MECHA can be helpful (Section 6). In closing, some concluding remarks and indications for future work are provided.

2. THE PROPOSED METHOD

Our work aims to bridge between a social view of collaborative activity and the work of designers of real systems who require systematic methods able to evaluate design choices. To this end, the method we propose supports the analysis and comparison of a set of design options. These differences are highlighted by describing scenarios that allow the analyst to focus on a specific case of use. The scenarios are introduced in the technical context of the current system. Subsequently, the scenarios are modified in the other options considered.

The reasons for this more qualitative approach is that evaluation of interactive systems is more economically carried out earlier in the development lifecycle, where redesign in response to identified problems is more feasible, as several authors (e.g., John and Kieras [1996] and Paternò [1999]) have pointed out. It is also worth noting that the scope of MECHA is broader than a number of other HCI evaluation techniques. For instance, many techniques in the HCI literature, such as heuristic evaluation [Nielsen 1993], focus on basic aspects of user interface design (such as speaking the user's language and simple and natural dialogues). MECHA additionally makes explicit reference to cooperation and hazards, thus allowing designers to carry out global evaluations of the impact of using different communication technologies in a complex safety-critical setting.

In addition, an inspection of the design options is carried out to make comparisons between them. The criteria for making comparisons are based on implications that design choices have for the tasks of individuals, for how they prevent or mitigate possible deviations in user behavior, and for the coordination of the activities of several collaborating individuals.

To capture the commonalities and differences between different design options, each of them is characterized by the specific media and representations that they provide for controllers, the specific ways that tasks are

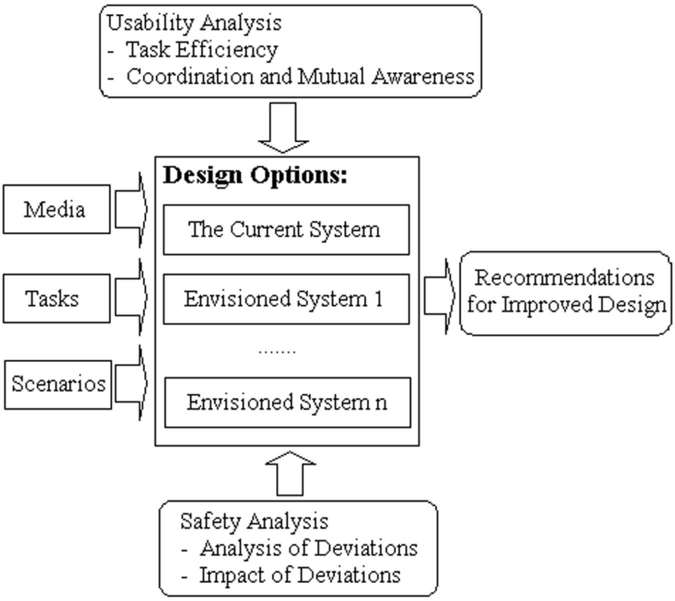

Fig. 1. The proposed method.

performed and allocated, and the scenario of use that is considered. The comparison of the design possibilities will be guided by a collection of criteria that address usability (for example, in terms of task efficiency, effective coordination, and mutual awareness) as well as safety (including analysis of user errors and their impact).

2.1 The Main Phases of the MECHA Method

A collection of criteria that can be used in making comparisons between the different design and task allocation options has been identified. The aim is not to provide specific measurable parameters that can distinguish in a quantitative way between the options, but instead to suggest criteria that form a framework in which designers can explore and analyze what the differences among the options are, thus understanding the design and the implications of design choices. Its results can be suitable to understand and represent the design rationale.

The method involves the following activities:

—*A task analysis and modeling of the current system and the identification of the possible main design options:* The approach we follow for task modeling is described in Paternò et al. [1998]. It allows designers to model cooperative applications and capture a wide variety of possible temporal relationships among the activities considered. However, MECHA can be used with other similar task-modeling techniques. The options differ (see Figure 1) in terms of the media available and how the tasks are performed and allocated, and on the choice of the artifacts and representations that are appropriate to support such tasks. In each option (including the current system) we create a matrix representing the balance of responsibilities between human and system.

—*Selection of specific sequences of tasks and elaboration of corresponding scenarios:* The purpose will be to create a context in terms of actions, technology, responsibilities, and environment, for the performance of the sequence of tasks identified, which depends on the design option being considered. In a safety-critical context, scenarios should be selected so as to focus on activities that can bring the system into a hazardous state.

—*An analysis of potential hazards and deviations belonging to a predefined set of categories:* For each design option (the current system and the envisioned ways to access new technology), the MECHA analysis of hazardous states is conducted using a method [Paternò et al. 1999], based on existing hazard analysis techniques [UK Ministry Of Defence 1996]. We ask questions that allow us to make hypotheses about some of the problems with the allocation of tasks and their performance that might arise. The result of this analysis can be summarized in tables indicating for each deviation the associated causes, consequences, protection, and recommendations. This yields a design rationale representation that is more suitable for interactive safety-critical applications than general-purpose techniques, such as QOC (Question-Options-Criteria) [MacLean et al. 1991], because it is more focused on analyzing the deviations and impact that they have on design.

—*An analysis of the coordination required among participants and to which extent the design supports those coordination requirements:* Many types of communication, more or less explicit, can occur in a work environment, such as a control room, and they must be captured precisely in order to provide effective support in an environment enhanced with new technologies.

2.2 Implications for Individual Tasks and Task Allocation

We can identify three main types of differences between the current system and "augmented" systems where new technology is available:

—*Change of task allocation between the human and the machine*: For example, user and interaction tasks that are performed routinely and require computations that are triggered by some automatically detectable event are good candidates for allocation to the system in more technologically evolved environments. In our case study of the data link environment, the update of the ground system (containing flight information) is no longer performed manually by the controller, but in an automatic way by the system.

—*Change of task allocation between human operators:* Old technology sometimes imposes constraints that can be overcome with more flexible technology. For example, both controllers working as a team can communicate with pilots by means of the data link functionality instead of only one controller having access to VHF radio communication, as is currently the case.

—*Change of objects manipulated by tasks and change of representations used to support tasks*: In this case, the introduction of new technology gives the possibility to use new artifacts to support the task performance. For example, in the new ATC systems the information formerly contained in paper flight progress strips can be electronically represented.

Furthermore, a number of factors relating to the way tasks are carried out must be considered when making comparisons between the design options. For instance, technological changes can have the effect of transforming interaction tasks into vigilance and monitoring tasks at which people are often less effective (cf. Hopkin [1988]). Similarly, design and task allocation decisions can have a significant impact on the workload of individuals and the range of responses to workload demands that are available to participants.

2.3 Hazards and Deviations

This collection of criteria is particularly important for interactive safety-critical systems, and involves studying the different failure and hazard characteristics of various design options. We use an inspection technique to go systematically through the actions that are required from participants, and consider ways in which failures might arise during a scenario, what the effect of failures might be, and what safeguards and defences exist in the system. Since an objective of the current work is to explore the impact of different arrangements of communication technology, a special emphasis will be placed on communicative actions.

The particular questions we will seek answers to are:

(1) What are the potential hazards that can arise as consequences of deviations, failures in communication, or erroneous actions in the scenario?

(2) Are there factors that tend to encourage miscommunication, erroneous actions, or faulty assessments?

(3) What recommendations concerning the user interface design can be provided to mitigate possible hazardous states and their effects?

This type of analysis will be performed with the help of *guidewords* (see UK Ministry Of Defence [1996], Leathley [1997], and Burns and Pitblado [1993] for related techniques). A *guideword* is a word or phrase that expresses and defines a specific type of *deviation*. Guidewords have been found to be a useful tool to stimulate discussion as part of an inspection process about possible *causes* and *consequences* in deviations of user interactions. Mechanisms that aid the *detection* or *indication* of any hazards are also examined, and the results are recorded. We have found it useful to investigate the deviations associated with the following guidewords:

—*None*, the task has not been performed, or it has been performed but has not produced any result.

—*Other than*, the task has been performed using the wrong data or producing wrong data.

—*Ill-timed*, the task has been performed at the wrong time.

In an analysis, these guidewords can be further refined. For example, *Other than* could be further refined into *Less*, *More*, or *Different*, indicating situations where the tasks have been performed using less, more, or different information. Likewise, *Ill-timed* can be refined into *Early* or *Late*, implying that the task is performed too early or too late.

The basic idea is that, for each design option, we consider the main tasks and the possible deviations that can occur in the performance of the task. Interpreting the guidewords in relation to a task allows the analyst to systematically generate ways the task could potentially go wrong, as a starting point for further discussion and investigation. Such analysis should generate suggestions as to how to guard against deviations, as well as recommendations about user interface designs, that might either reduce the likelihood of the deviation or support detection and recovery.

We believe that this type of analysis is particularly effective when performed by a multidisciplinary team, as in our project, that includes designers, software developers, and end-users. Each contribution can enrich the discussion with specific expertise and knowledge, helping to identify possible deviations and the most effective design choices.

2.4 Coordination of Activities and Mutual Awareness

The concept of "articulation work" and the means by which the activities of individuals are coordinated are a complex topic. For current purposes, we focus on one aspect, namely, the way in which technology changes (such as the introduction of data link) have an impact on the kinds of coordination that are necessary and possible. More specifically, the two questions we will be asking about the design alternatives are:

(1) What coordinations are needed so that the tasks of the users, such as two controllers, are brought into step? The answer to this question will typically be dependent on the particular roles, responsibilities, and tasks of the individuals involved.

(2) How such coordinations will be supported by the available mechanisms? The answers to this question are likely to be dependent on the details of the technologies and artifacts that mediate the tasks of individuals and communications between them.

3. A CASE STUDY IN AIR TRAFFIC CONTROL

3.1 The ATC System

The principal objectives of the air traffic control system are generally stated as the achievement of safe and expeditious flow of traffic through airspace. In this case, safety is interpreted as meaning that separations standards for the minimum safe distance between aircraft should be respected. The goal of expedition means that, so long as safety is preserved, the flow of traffic through the airspace should be maximized.

Civil airspace is partitioned into a number of regions known as *sectors*, by horizontal and vertical divisions. In a typical air traffic control center, aircraft within a sector are managed by two air traffic controllers, who work closely together, but individually have rather different roles and concerns.[1]

The *executive* controller is able to contact aircraft using VHF radio and is directly responsible for making short-term decisions and maintaining the appropriate separation distance between aircraft. The *strategic* controller, on the other hand, has longer-term concerns and coordinates with strategic controllers of adjacent sectors to agree on flight parameters (particularly the altitude) of aircraft entering and leaving the sector, so that passage between sectors happens in a safe and orderly manner.

3.2 The Options Considered

In our case study, initially three design options were analyzed and compared. These were the current ATC system, and other two options that differ in the way that new communication technology is used. The two new options both use data link (in different ways) and replace the paper strips with other electronic artifacts. This has an impact on how the tasks are performed and allocated, and on the choice of the representations that are appropriate to support such tasks.

It is useful to explain why we selected these options for evaluation. We wanted to consider a small number of options that allowed us to address the main design issues. The key distinction between the new variants is that one aims to replicate existing communication patterns and support them with the new data link technology in order to overcome the reported bottleneck of overloaded radio channels. So, it permits an investigation of the properties of the communication medium, without altering the division of labor between controllers. The other one aims, instead, to change the allocation of communication tasks among controllers, taking advantage of the possibilities of the new technology and using different software artifacts to support controllers' tasks. It investigates whether it is possible to optimize the allocation of the data link to both controllers with a different allocation of tasks.

[1] This study looked specifically at en-route air traffic control in France, but much that is said here about that system and the technical innovations that are taking place within it also applies to ATC environments elsewhere.

However, for sake of brevity, the analysis of only two options will be presented here, as the purpose of this article is to explain how the method works rather than to provide details of the case study. The options selected for detailed consideration in the article are the current system and an envisioned system with data link communication available to both controllers where the task of managing aircraft when they change sector becomes the sole responsibility of the strategic controller, and is connected with the task of negotiating transfer parameters.

3.3 The Aspects Considered

For sake of brevity in the article we consider only one scenario. The purpose of using a scenario will be to enable the analyst to envision how a particular technological configuration will affect the attempt to solve specific problems in a specific context. The same scenario is used to consider each of the design possibilities. First, it is introduced in the current ATC system. Later, the differences for the option discussed will be described.

We structure the description of scenarios around a “template” that has been used previously as part of a scenario-based error analysis technique [Fields et al. 1997]. The template provides sections for describing the *agents* who play a part in the scenario, the physical *situation*, the physical *environment* that generates problems for the agents, the *tasks* to be carried out, *systems* and devices that are present, and the *actions* that take place as the scenario unfolds.

In the selection of the tasks to consider we use a level of “task granularity” chosen pragmatically to allow us to discuss and reason about the pros and cons of each arrangement. The tasks considered in the current ATC system are as follows:

—*Detect Problem:* Either controller *can identify* a possible conflict in the current air traffic.

—*Inform Executive:* Strategic informs executive that a problem has been detected or that something has to be done.

—*Handle First Contact:* The executive controller responds to the first communication from a pilot who has just entered the sector.

—*Solve Problem:* The executive controller’s cognitive task of finding a solution to solve a conflict or to achieve a more expeditious flow of traffic.

—*Send Clearance:* The executive controller’s task of sending instructions to aircraft.

—*Update Strip:* The executive controller annotates the paper flight progress strips to record the status and history of traffic in the sector.

—*Update Ground System*: Both controllers, though usually the strategic, are able to update data in the ground-based computer system.

—*Negotiate Transfer Parameters*: The strategic controller negotiates with counterparts of adjacent sectors in order to agree on parameters for aircraft moving from one sector to another.

—*Monitor Radar:* Both controllers, especially the executive, monitor information provided by the radar screen.

—*Handle Last Contact:* The executive controller instructs the pilots of aircraft leaving the sector to contact the executive controller of the new sector on an appropriate VHF radio frequency.

4. THE CURRENT SYSTEM OPTION

4.1 The Systems and Its Usage

4.1.1 *The System.* Looking at the environment of a single working position (see Figure 2), we see a complex array of devices and information artifacts. Controllers are provided with radar screens, on which the location and parameters of aircraft, the airways, and the sector boundaries are displayed. A VHF radio is provided for use by the executive controller (seated on the left in the figure) to communicate with pilots in the sector (this possibility is available also for the strategic only for emergency cases). Telephones allow the strategic controller (seated on the right) to keep in touch with other strategic controllers of neighboring sectors. A Touch Input Device (TID) allows the ground-based computer system to be updated to reflect changes to flight data (such as the time, flight level, and track). Usually, it is the strategic controller who performs these updates.

Thus, in the current system three kinds of communication exist:

—Between strategic and executive controllers of the same sector (for example, voice and "elbow" communications to attract attention).

—Between the strategic controllers of adjacent sectors concerning aircraft passing from one sector to the next (carried out using the telephone).

—Between the executive controller and pilots of aircraft in the related sector (by means of VHF radio).

The controllers on a working position have also some flight paper strips printed "in real time" and delivered to a given control position about 10 minutes before the flight enters the sector handled by that position. There is one strip per aircraft: the strip gives general information on the aircraft (identifier or callsign, aircraft type; see Figure 3) but also flight information on the planned route of the aircraft. All the data printed on a flight strip are manually "updated" by either controller using a pen.

4.1.2 *The Scenario.* The scenario is centered on the activity of a single sector, and therefore an executive and a strategic controller are involved. Three other agents are also important: the pilots of aircraft in the sector (in particular, BAW1234 and AZA5678) and the strategic controller of an

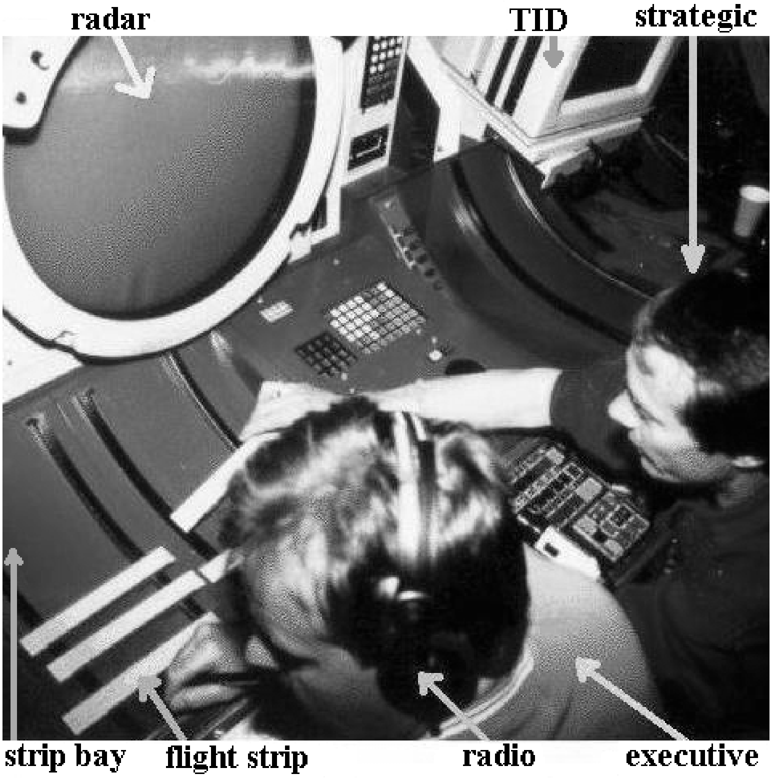

Fig. 2. The current air traffic control position.

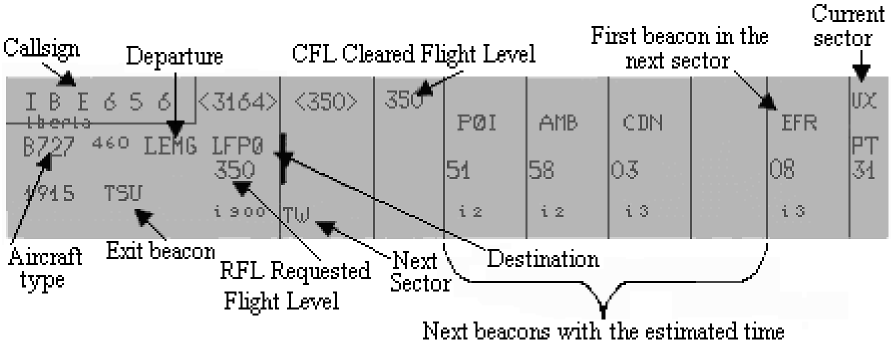

Fig. 3. Example of paper strip.

adjacent sector. These five agents will be referred to as **strategic**, **executive**, **BAW1234**, **AZA5678**, and **adjacent strategic** respectively.

4.1.2.1 *Situation and Environment*. The activity takes place in an en-route sector in which two airways, A and B, intersect. A third airway also intersects A. The sector has boundaries with other en-route sectors. The two aircraft in question, BAW1234 and AZA5678, are flying along routes as

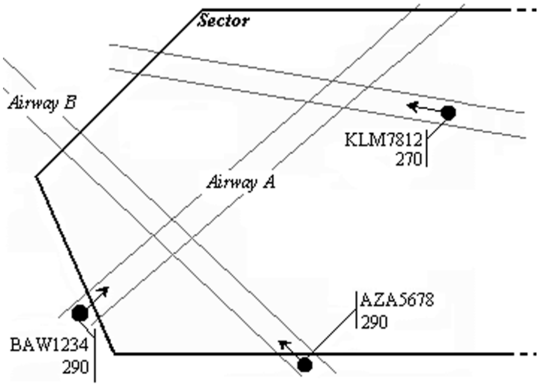

Fig. 4. Intersecting airways in an en-route sector.

shown in Figure 4. Both aircraft are at Flight Level 290 (FL 290). The problem faced by the controllers is the potential conflict between these two aircraft at the intersection of the two airways.

A third flight, KLM7812, may generate a conflict depending on the solution to the first conflict. We assume that the executive is unable to make the aircraft change direction or speed to solve the conflict; and instead, vertical separation will be achieved.

4.1.2.2 *Task Context*. In this scenario, the conflict between the two aircraft is detected by the strategic controller. A strategy for avoiding the conflict (by altering the altitude of one aircraft) is devised by the executive, who issues a new clearance to the relevant aircraft. The strategic controller is responsible for seeking agreement for a revised coordination with the appropriate adjacent sectors.

4.1.2.3 *System Context*. In this option we assume a technological context similar to that found in ATC centers today.

4.1.2.4 *Actions*. After having fixed the transfer parameters of the two flights with the strategic controllers of previous and next sectors, the strategic controller detects the potential for conflict using flight progress strips. This involves comparing flight levels and estimated times of arrival for the two aircraft at the point at which the airways intersect. At this stage, both aircraft are still outside the sector considered (Figure 4). The strategic controller then informs the executive (who will devise a strategy for avoiding conflict) and at the same time moves the strip into the correct position in the strip bay, within reach of the executive.

After few minutes, BAW1234 enters the sector, and the first contact is received from the pilot of BAW1234 via the VHF radio ("Bordeaux control, Speedbird1234, level 290, good morning"). Similarly, AZA5678 announces its arrival into the sector. On the basis of information previously supplied by the strategic controller, the executive is able to decide how to avoid the conflict. In this particular scenario, the executive decides to solve the

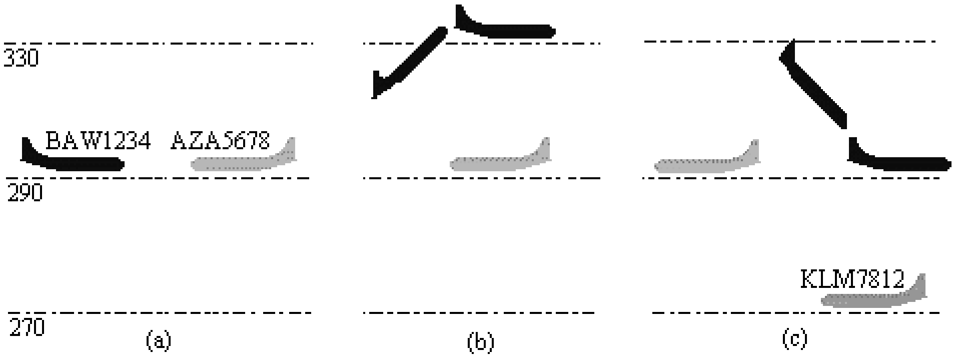

Fig. 5. The different flight phases in the scenario.

conflict by increasing the altitude of BAW1234 (see Figure 5(a)). This vertical separation will also avoid a conflict with flight KLM7812 that is traveling on the third airway. Therefore, the executive controller issues the clearance, which is then acknowledged by the pilot in question:

Executive: Speedbird1234 CLIMB TO flight level 330

BAW1234: Speedbird1234 CLIMBING TO flight level 330

The executive then annotates the paper strip to indicate that the clearance has been given and that the aircraft is climbing (see Figure 5(b)), and starts to monitor on the radar if the flight actually performs the instruction. As this instruction affects the previously agreed Transfer Flight Level (TFL), at which the aircraft will make the transition between the current sector and the next one, some renegotiation will be needed. The strategic controller, on hearing the communication between the executive and pilot, recognizes the need to contact the neighboring strategic controller to confirm whether the revised transfer flight level is acceptable.

The adjacent strategic controller, however, is unable to accept the revised altitude of BAW1234 as the TFL. An agreement between strategic controllers is reached on the original TFL: 290. The strategic controller now informs the executive verbally, and subsequently updates the flight progress strip to reflect this change.

Therefore, another exchange between the executive and pilot will be needed to transmit the new transfer flight level to the pilot of BAW1234. This must occur sufficiently far in advance to allow the aircraft to make the descent before changing sector, and crucially, it must occur after BAW1234 has crossed airway B (see Figure 5(c)):

Executive: Speedbird1234 DESCEND TO flight level 290

BAW1234: Speedbird1234 DESCENDING TO flight level 290

The executive then updates the strip, and one of the controllers updates the ground system (so that an updated flight progress strip can be printed in

Table I. Scenario Actions for Option 1

1 Strategic determines the possibility of conflict and alerts the executive.
2 BAW1234 enters the sector with the VHF message: "Bordeaux control, Speedbird1234, level 290, good morning"
3 Executive solves conflict detected by strategic issuing the clearance "CLIMB TO 330"
4 Positive answer received from BAW1234 signalling that BAW1234 starts to climb. Paper strips and ground system annotated manually
5 Strategic overhears communication and recognizes need for renegotiation of transfer flight level with the adjacent controller
6 Strategic contacts adjacent strategic controller
7 Agreement reached on TFL 290
8 The strategic informs executive that the executive has to send a new clearance for the TFL 290
9 The executive sends the clearance "DESCEND TO 290"
10 Affirmative answer received from BAW1234. Paper strips and ground system annotated manually
11 Last contact from executive: "Speedbird1234 CONTACT Marseille control 135.85, GOODBYE" followed by pilot's reply
12 Paper strip removed

the next sector). Afterward, the executive performs the "last contact" by informing BAW1234 of the frequency of the next sector. Once BAW1234 is no longer under the control of the current executive, the executive can remove the associated flight progress strip from the bay and discard it.

4.2 EVALUATING THE DESIGN

4.2.1 *Task Performance.* Table II shows how tasks are allocated between agents in the current ATC system. In this option we can note that (1) there is no computerized support from the system for performing the main tasks identified and (2) that there are various tasks that can be performed by both controllers. Similar tables will be provided for the other options to highlight how tasks and their allocation are modified within them (for example, in some cases some functions will be carried out by the system). Some of the differences will arise as a consequence of the introduction of data link. For example, in the next option we will have a "change flight data" task instead of "update strip" because the next option is a stripless environment.

The main limitations of the current system are especially highlighted in situations of high traffic. Traffic increases (and consequently the increase of conflicts to be prevented, detected, and solved) are managed and controlled largely by the executive, although the strategic constantly pays attention to monitor both the executive's activity and the traffic situation.

There are problems in situations of very high traffic because the executive remains the only agent in charge of making problem-solving decisions and communicating them by means of radio communications without automatic decision support. Therefore, the greater the number of aircraft in the sector, the harder the executive's task, and the bigger his or her

Table II. Task Allocation in Option 1

Strategic	Executive	System
Monitor Radar	Monitor Radar	
Negotiate Transfer Parameters		
Update Strip	Update Strip	
Update Ground System	Update Ground System	
Detect Problem	Detect Problem	
	Solve Problem	
	Send VHF Clearances	
	Handle First Contact	
	Handle VHF Last Contact	
Inform Executive		

workload is in coordinating all the activities needed to ensure traffic safety and regularity. With a bigger workload, executives are much more prone to introduce errors and omissions in performing their tasks. In addition, congested radio channels increase the possibility of misunderstanding due to simultaneous communications requiring the communication to be repeated, and in the worst case, force the speaker to wait until the frequency is available again, making the task of communicating a message more difficult.

If we analyze the type of controllers' tasks, we note that they are different in *number of tasks*: the executive has to support all the communications with pilots and, at the same time, must respond to and solve problems as they appear in real time. The strategic controller, apart from the task of monitoring the system, has only to negotiate with other strategic controllers the flights' transfer parameters and to update the ground system for all the vocal communications that occur between the executive and pilots currently crossing the sector.

In addition, the executive's work is more demanding, because of *time constraints and deadlines*. The executive must often solve problems and make decisions before problems develop, whereas the temporal requirements on when the strategic controller updates the ground system, and so on, can be rather less stringent. Therefore, in contrast to the executive controller, the strategic controller can "organize" the work with a degree of flexibility.

We also note that controllers' tasks are different in regard to the *type of skill requested*, because for example the task of resolving an unforeseeable conflict quickly in the traffic flow is obviously more demanding compared to the strategic controller's work of updating the ground system (that is a "routine" task above all).

The above considerations highlight that the executive controller is engaged in supporting heavier activity than the strategic (in terms of tasks, constraints on them, and skill requested) which results in an imbalance in the allocation of work between the two controllers.

Regarding the representation of data in the current system, the primary sources of information are the radar and the paper flight progress strips.

When we analyze such strips we should not concentrate only on the information contained in the strips, but also in the type of interaction that a specific representational form allows supporting the users' tasks. For example, paper flight progress strips are generally considered by controllers to be an extremely flexible tool, supporting both visual and tactile memory (see Bentley et al. [1992] and Mackay et al. [1998]). Additionally, the ability to easily rearrange and reorder strips in the strip bay (see Figure 2) depending on different criteria can play a crucial role in conflict detection and decision making. Besides, recall that the strip offers a useful means of communication between the strategic and the executive controllers. The two controllers can work simultaneously on the strip bay, annotating, moving, and pointing at strips (for example, the particular position where the strategic puts a strip in the strip bay is used to communicate a different level of urgency with which the executive has to put attention to it).

In addition, on the one hand the task of writing on paper strips can be an important means by which the controller's mental "picture" of the traffic situation is updated and reinforced. On the other hand, however, strips are sometimes seen simply as a distraction from looking at the radar screen.

The main controllers' task of avoiding, detecting, and resolving conflicts results from a repeated cycle of looking at the aircraft representation on the radar and the associated flight strip in the strip bay. Thus, in the current system, especially in high traffic situations, the continuous activity of doing it (moving the eyes from two different artifacts—the radar and strip bay—and coordinating the information from these sources) should not be overlooked in any assessment of controller workload.

4.2.2 *Hazards and Deviations*. The result of our analysis can be represented in tables such as Table III, where for each deviation, designers can indicate causes, consequences, protection, and recommendations. This gives a good documentation for design rationale more suitable for interactive safety-critical applications than that provided by approaches such as QOC [MacLean et al. 1991] because it is more focused on those aspects that can have an impact on the safety of the environment considered.

In the case of giving a clearance in a VHF-only environment (for example, "Speedbird1234 climb to flight level 330") if we consider the "*none*" deviation it can indicate either that the controller has not sent the clearance or that it was sent but not received.

If the executive does not send a clearance there can be a number of possible explanations. As we saw in the scenario description, clearances can be the result of various activities: a conflict is detected, communicated to the executive, and then a solution is identified by the executive and translated into possible clearances. Each of these phases has the potential to fail. For example, the possible conflict can remain undetected, or it may be detected; but the strategic is interrupted before informing the executive, and the conflict is subsequently forgotten.

Table III. Example of Analysis of Deviations Based on Guidewords

Task: Handle VHF Last Contact				Guideword: None
Deviation	Causes	Consequences	Protections	Recommendations
No message is sent by the controller	Controller fails to recognize the need for giving a new frequency (e.g. mistake due to failing to note aircraft location in proximity to sector boundary) **Perception Problem** Controller recognizes need, but fails to send message (e.g., memory lapse or high workload-induced slip) **Action Problem**	No new frequency received by pilot—aircraft may enter the sector without having contacted the new controller	(1) If aircraft enters next sector without making a call, the next controller will call the current one (2) Pilot calls ATC	Provide an indication when aircraft within threshold distance from sector boundary. Such an indication would persist until electronic system updated
No message is received	Pilot fails to perceive message (e.g., inattention, high workload) **Perception Problem** Total failure of communication technology		(1) Controller expects answer; lack of it may prompt a second call (2) Pilot calls ATC	If there is no answer within some time limit an alarm message (e.g., an audible signal) could be automatically activated

In this option the problem of the pilot failing to hear the clearance could arise if other tasks on the flightdeck distract the pilots. It may take some time before the executive realizes that the communication has failed. As the reader can see this analysis of deviations can be represented by tables structured as that below that give the analyst or designer a means of recording the justification or rationale for decisions that have been made. In the Table III we consider the task of sending the "last contact," instructing a pilot leaving the sector to contact the executive of the next sector. Each cause of deviation can be associated with the phase of the interaction cycle (intention, action, perception, understanding) when it occurs.

An *"other than"* deviation with the voice communication can be interpreted as meaning that, for some reason, the executive either sends the wrong information (wrong clearance, wrong parameter, or both), or sends right information but it is incorrectly perceived by the pilot, or the wrong pilot is contacted.

Wrong information can be given for various reasons. For example, a relevant environmental factor may have not been taken into account. In this scenario, a third aircraft, KLM7812, is also in the sector, and the solution proposed can generate another conflict later on. Also, the *type* of solution may be correct, but some of its details may not be. For example, the flight level chosen may not be appropriate for the direction of travel.

Similar problems in a VHF-only environment can also happen because there is a lack of automatic tool support for making decisions. Thus all the support available for controllers—paper strips and radar information—require considerable cognitive effort. The pilot may misperceive or misinterpret the controller's command for various reasons, both linguistic and contextual.

Ill-timed communication is one that occurs either too early or too late. If the executive communicates a solution too early he or she can generate new problems other than solving the current one. For example, changing the flight level too early can generate a conflict with a flight that originally was not in conflict with that under consideration. Various reasons exist for a communication being late; most obviously, either controller may be busy with other duties. Another possibility is that when the executive needs to communicate, the radio channel is already "occupied" by another pilot. Indeed, another problem can be derived from limitations of radio communication. Since only one speaker can broadcast over the frequency at a time, the resulting communication has an asymmetrical nature (*one* speaker, *many* hearers). As all the pilots compete for the use of this resource, such sharing can be seen as penalizing pilots.

4.2.3 *Mutual Awareness and Cooperations.* In the scenario we are considering here there are three main points at which the streams of activity in which the executive and strategic controllers are engaged in must be brought together and coordinated:

—The first coordination occurs when the strategic controller determines that a conflict is likely to occur, and informs the executive controller (step 1 in the scenario). This is a very explicit form of coordination between the two streams of activity.

—The second point of coordination is that the strategic controller must know the conflict-avoiding clearance the executive has issued ("*climb to flight level 330*"—steps 3, 4, and 5 in the scenario) so as to know that negotiation with the strategic of the adjacent sector is necessary.

—The third point of contact is that the executive must know the outcome of the negotiation between the two strategic controllers, so as to be able to pass the appropriate clearance to the aircraft (step 8).

Having identified what coordinations are required for the tasks on individuals to be carried out successfully, we can now begin to look at how they take place. In other words, what "coordination protocols" [Schmidt and Simone 1996] exist, and how they are achieved?

In this situation, where VHF radio is the only form of communication between the ground and the air, communications between pilots and the executive are public, in the sense that they are audible in the control room, and may be overheard by the strategic controller. The communication between the strategic controller and controllers of adjacent sectors, on the other hand, is conducted through a medium with different properties: the telephone. One property is that conversations conducted using it are private, and the executive is not party to them (at best, only one half of the conversation is made public in the control room).

The second coordination can therefore be accomplished in two ways: firstly, by the strategic controller monitoring the talk on the VHF channel and responding when necessary, or secondly, by explicit action on the part of the executive. This explicit notification can take place in a number of ways (speaking, writing on flight progress strips, nudging, and pointing). In practice, all these mechanisms are used, and the choice depends on a number of factors, such as the level of ambient noise in the control room, and the other tasks that are being carried out concurrently.

The third coordination is more constrained in how it can take place: since the communications between sectors cannot be overheard, in the case where a renegotiation of a new transfer flight level occurs, the strategic controller must inform the executive explicitly. On first sight, this constraint on the third sort of coordination may seem like a deficit, a lack of flexibility of coordinative mechanisms. However, this asymmetric arrangement (where, by monitoring, the strategic controller can gain an awareness of the tasks of the executive, but not vice versa) is entirely consistent with the idea that a “protective cocoon” is constructed within which the executive controller works [Hughes et al. 1992].

The status of the executive’s work is rendered public by the nature of the VHF channel. This will tend to have an impact on the nature of the division of labor that exists between the controllers. The strategic controller is able to “help out” in times of busyness (a clear connection between coordination and safety). Flexible arrangements such as this are made possible by the public nature of some of the communication media. Our own observations at a current air traffic control center have identified a number of routine situations where the “shared space” of VHF transmissions permits just such an arrangement. The strategic controller is able to appropriate some of the tasks normally carried out by the executive (for instance, those connected with flight progress strip management).

5. THE “DATA LINK FOR BOTH CONTROLLERS” OPTION

5.1 The System and Its Usage

Here we consider an environment with more electronic support. We thus consider the possibility that both controllers, in different circumstances, can send clearances by data link, and they interact with enriched flight labels instead of paper or electronic strips.

5.1.1 *The System.* In this option the strategic controller can use data link technology in the flight transfer phase. This means both frequency-related clearances and clearances related to change of flight parameters during such a phase. The interactions with system functionality are supported through the radar labels, as the ability to operate directly on an item in the current focus of attention is really important for the controller:

in some sense the radar data block becomes a multiline label that replaces the flight progress strip. Note that the information permanently displayed on these labels is kept to a minimum, in order not to clutter the screen, allowing at the same time the controller to expand the label (moving and clicking the mouse onto it) when he or she has to up-link some instructions or collect aircraft information.

So, in this improved system the flight parameters of aircraft currently in the sector are displayed "on demand" by controller when necessary, leaving on the screen only a minimal kernel of the most relevant information (the "standard" label; see Figure 6(a)). In Figure 6 we show an example of session when the controller wants to send the instruction of changing the flight level parameter to a pilot. Moving the cursor onto the radar label, the label is expanded, and it is possible to select the CFL instruction, the requested new flight level (330 in our example; see Figure 6(b)), and then either to abort or to send it. When the clearance is sent to the aircraft the symbol beside the callsign has a special color (orange) to indicate that the Wilco (the positive answer from the pilot) has not been received yet (see Figure 6(c)). When the "Wilco" is received (Figure 6(d2)), the symbol returns to the normal color, and the icon of the aircraft changes to indicate that it is actually climbing. Another possibility is that the pilot is unable to perform the instruction: in this case, an error symbol beside the callsign is displayed in a warning color (Figure 6(d1)). In such a way, the improvement is that for each aircraft currently in the sector, the attention of controllers is focused only to its (multiline) label, which allows them to perform their activities, to send clearances to pilots, to collect flights parameters, etc., and the attention is no longer divided between the representations of aircraft on radar and the associated strips.

5.1.2 *The Scenario.* With respect to the previous option, the difference is in terms of actions. Suppose that the strategic controller has to support all the activities necessary during the change-of-sector phase, leaving the normal en-route phase to the executive. Given this option, it becomes crucial to precisely define when the data link control passes from the strategic to executive, and vice versa, in order to make it clear when a flight is under the control of one or the other. In our scenario, we assume that the first clearance (climb) has to be sent when the BAW1234 is into its en-route phase, and the second clearance (descend) has to be sent while BAW1234 is into its change-sector phase.

5.2 Evaluating the Design

5.2.1 *Task Performance.* The workload between controllers is more distributed, both in terms of number of tasks and in terms of type of task. Under this option the strategic controller has a more direct role in making decisions and controlling the traffic as far as it concerns the flights in their changing-sector phase, resulting in greater involvement in the overall controllers' activity of managing and controlling the traffic. As shown in

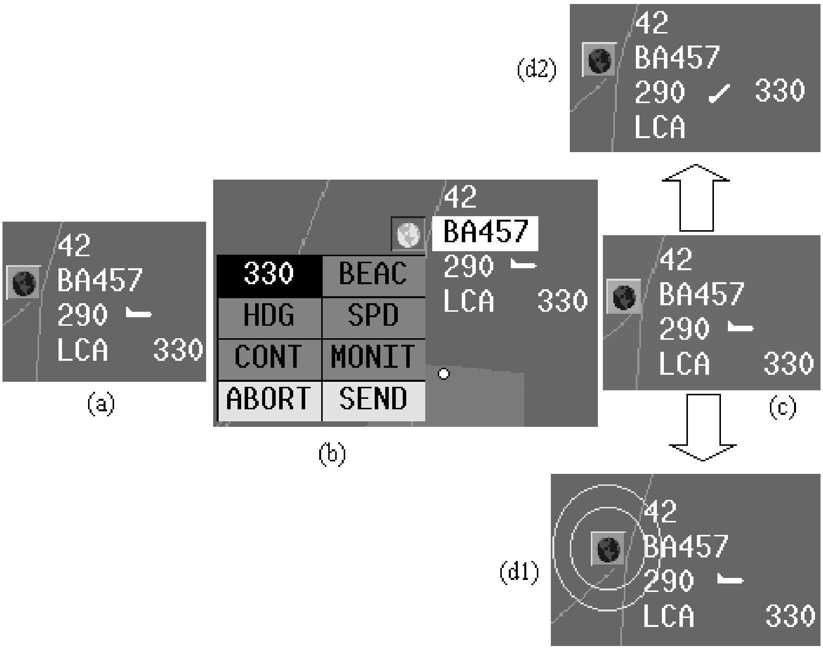

Fig. 6. An enriched flight label.

Table IV. Scenario Actions for Option 2

1	Strategic determines the possibility of conflict and alerts the executive
2	The ground receives the data link message when BAW1234 enters into the sector: "MONITORING 120.5"
3	Executive solves conflict detected by strategic sending the data link clearance (CLIMB TO 330)
4	WILCO received from BAW1234 => ground system automatically updated
5	Executive makes strategic aware of need of renegotiation on TFL
6	Strategic contacts adjacent strategic controller
7	Agreement reached on TFL 290
8	The strategic sends the data link clearance (DESCEND TO 290)
9	WILCO received via data link from BAW1234 => enriched flight labels and ground system automatically updated (change flight data)
10	Last contact from strategic (MONITOR Marseille 135.85) to BAW1234 and pilot's reply (via data link)

Table V the task *Update strip* disappears and is replaced with a *Change flight data* task, as there are no more strips in this environment.

In fact, whereas in the other cases the strategic controllers had to perform all "routine" activities, now their activity can have a direct impact on the general system speed-up, as all flights going into the sector or leaving the sector have to communicate with them.

On the one hand, this can result in less work for the executive, and in true "parallelism" by allowing clearances to be sent to pilots (in different flights phases) by both controllers, overcoming most limitations of the

Table V. Task Allocation in the New Option

Strategic	Executive	System
Monitor Radar	Monitor Radar	
Negotiate Transfer Parameters		
Change flight data	Change flight data	Change flight data
		Update ground system
Detect Problem	Detect Problem	
Solve Problem	Solve Problem	
	Send VHF Clearance to Pilot	
Send DL Clearance to Pilot (only during change sector)	Send DL Clearance to Pilot (only during en route)	
	Handle VHF First Contact	
Monitoring Frequency		
Handle DL Last Contact		
	Handle VHF Last Contact	
Inform Executive		

previous options. On the other hand, the need for mutual awareness and coordination between the controllers becomes more acute.

As a matter of fact, a “stripless” environment actually gives controllers more time to watch the display, even though controllers have to spend some time for the additional activities related to mouse selection of enriched labels to display the data. The problem of linking the aircraft representation with the information normally included in the flight strip is addressed by this design. However, this raises additional questions about how well particular tasks (for example, the task of comparing values associated to two or more aircraft in order to avoid conflicts) could be performed under this option. One possibility is the use of an additional screen with electronic strips interactively supporting such tasks. Further improvements could be imagined: for example, the possibility to send specific clearances in a graphical manner rather than in a textual way, or selecting on the screen the next beacons rather than textually specifying them.

5.2.2 *Hazards and Deviations.* In this option possible deviations with respect to the expected behavior can be originated when the flight is in a position very close to the change sector phase so as both controllers can believe that it is under the other controller’s responsibility. This requires a user interface mechanism highlighting in which of the two phases the flight is (for example, adding an internal border to the sector indicating when the flight enters in the area requiring change sector clearances).

Having a controller dedicated to handle the change sector phase with data link commands decreases the possibility that the related clearances are sent too early or too late (or they are not sent at all). However, there may still be some problems when several flights are transferring, and one of them requires cross-sector negotiation by the strategic controller, thus distracting the strategic from sending the clearances concerning the other flights.

In the *Other than* deviation case we can imagine the possibility that the executive identifies a good logical solution, but it supplies incorrect parameters. For example, the executive may attempt to use a flight level that is reserved for flights traveling in the opposite direction. It is easier in a graphical user interface to include some support for avoiding such situations, for example, by disabling the possibility of assigning a flight to an altitude associated with the opposite direction. This can be done by exploiting the more electronic support foreseen in this option.

5.2.3 *Mutual Awareness and Cooperation*. Many of the coordination issues of the previous case are relevant here: the executive makes tactical decisions (of which the strategic controller should be aware), and the strategic controller manages cross-sector negotiations (whose results the executive needs to know).

Several observations can be made about the mechanisms by which coordination may be achieved in this case. In contrast to the VHF-only case, there are fewer “public” mechanisms to support coordination. However, the more sophisticated environment does contain some cues for coordination. For instance, the fact that an aircraft has been assumed under the control of the executive of the new sector is indicated directly in the display (the color of the aircraft symbol), thus providing a kind of “computational coordination mechanism” (to use the terminology of Schmidt and Simone [1996]).

This arrangement of communication media has a number of implications for the coordinations that must take place. The first and third coordinations are relatively unchanged from the VHF case, and still require an explicit action on the part of the strategic controller to notify the executive. The second point of coordination is different in this case, compared to the previous one. It is no longer possible that the activities of the two controllers can be coordinated by the strategic controller “listening in” to talk on the VHF channel. Some more explicit means of communicating will be necessary (such as the executive notifying the strategic controller when the clearance has been issued).

One further implication of this is that the technology may tend to limit the kind of opportunistic intervention and “helping out” that is made possible in the current system, through shared media (such as the VHF channel) that allow controllers to gain an awareness on one another’s work. An implication for design is that if such flexible working is deemed desirable, then additional measures will have to be taken to permit the levels of mutual awareness that are needed to support it, by compensating for what the replacement of voice by data link has tended to diminish.

The most significant change from the previous case is that the strategic controller now carries out a task that was previously carried out by the executive (arranging for the transfer of the aircraft to the next sector). This might have the effect of changing the executive’s workload and improving the balance of work between the two controllers. It also changes coordination requirements: (1) it is no longer required that the strategic communicates the change sector flight parameters to the executive and (2) the executive controller must become aware of when the aircraft is no longer under his or her control because the aircraft is in the area under the control of the strategic. Several mechanisms exist to facilitate this kind of coordination, the most obvious being the color coding used to indicate the status of aircraft.

6. SUMMARY AND LESSONS LEARNED

In our study we have seen how the MECHA method can be used to analyze and evaluate the impact of introducing a new technology in a safety-critical context. The method involves a consideration of task allocation and performance, deviations, and cooperation mechanisms, in a way that can be useful in many application domains. However, some aspects, such as the analysis of deviations and suggestions for preventing them or mitigating their effects, are particularly important in safety-critical systems, where the effects of actions often cannot be undone and may be catastrophic.

In addition, we have seen how safety-critical applications are interactive real-time applications where users (controllers in our case study) should be ready to handle unexpected situations. As controllers can perform a number of activities, in some cases by cooperating with other controllers, various types of deviations can occur during their performance. In these possible concurrent activities the root of many safety issues lies. Such a concurrency often allows more efficient and flexible performance of tasks, so the problem is not to limit it but to design environments able to effectively and safely support it.

Table VI summarizes some of the findings for each of the two design options that have been discussed in the previous sections.

We have seen how a combined analysis of the tasks carried out by the user and the possible deviations provides information useful for designing applications able to improve task performance and safety. This analysis cannot leave out of consideration the specific environment (in terms of single or multiple agents, different media and artifacts, object representations, etc.) where the tasks are performed. In this sense a comparative analysis among current and envisaged systems can give useful information to highlight for example which arrangement allows a better performance of which tasks (hardly ever the case that one system is the "best" in any absolute sense). In our analysis we have also indicated a useful lesson, which is to design artifacts that allow *users to better integrate information concerning a specific logical object:* while paper strips require a strong effort to integrate information that is on them with that provided from the radar, enriched flight labels decrease such an effort by supporting integrated information directly on the radar screen.

Although increased automation does not always improve usability [Bainbridge 1983], the introduction of more electronic support can lead to a *decrease in the workload* of controllers because of the possible automation of some tasks (for example, whenever a command is accepted by the pilot

Table VI. Summary of Design Options

	Task Performance	Hazards and Deviations	Mutual Awareness and Cooperation Mechanisms
Only VHF	Low with high traffic (mainly due to combined bottleneck of VHF media available only to executive)	Considerable (due to executive bottleneck)	*From executive to strategic:* implicit (because of VHF communications with pilots) *From strategic to executive:* always explicit (for phone communications with other strategic)
D.L. for both	Improvements of global performance: two media for communication with pilots—VHF and DL—available to *both controllers*	Low (if both controllers are aware of when each flight is under the control of which controller)	*From executive to strategic and from strategic to executive:* Need for explicit mechanisms for all data link communications with pilots

the system is automatically updated). It is also possible to introduce functionality to automatically identify information useful for *supporting decision-making* activity.

The systematic analysis of deviations gives useful suggestions for improving the design. For example, it is possible to identify when an action from the controllers is required, and to suggest introduction of warning messages to capture their attention. The level of intensity of the alarm messages should be related to the kind and severity of hazards that may arise. Such messages should aim to *reduce the likelihood that deviations from expected behavior may have an impact on the safety of the system*.

In addition, we have seen how our method is able to understand the established practices of controllers that have produced *a set of coordination mechanisms that are well recognized and sufficiently reliable*. This is important because this level of coordination and robustness should be preserved in future and possibly even improved.

The MECHA method requires a considerable effort from the evaluators because it requires that various activities be performed: task analysis, identification of design options, analysis of deviations, and cooperations in the context of a set of meaningful scenarios. However, we are considering applications where it is extremely important to guard against actions that can have a high impact, pose a threat to human life, and whose negative effects cannot be undone. Thus, it can be justifiable to spend more time in the design and evaluation phases if this allows evaluators to reach a better understanding of the possible consequences of user deviations and obtain a design that can either prevent them or mitigate their safety-critical effects when they occur.

The resulting approach allows designers to bridge between a notion of collaborative activity and the work of designers of real systems who require systematic methods able to evaluate design choices thus overcoming limitations of alternative approaches. For example, ethnographic-based approaches tend to provide many details useful in the analysis phase but less support in the design whereas traditional task analysis has paid little attention to the impact of possible user deviations on design.

7. CONCLUSIONS AND FUTURE WORK

The design of user interfaces is a complex process that must take into account many different factors. This is especially the case when designers consider a cooperative and interactive safety-critical application, such as the air traffic control example considered in this article. When such safety-critical applications are analyzed, both usability and safety have to be carefully considered in an integrated way.

We have presented a method based on the use of three types of criteria (implications for task performance and allocation, analysis of deviations, and coordination) and its application to en-route air traffic control, by considering some design options in the use of communication media. This analysis has shown how it is possible to improve the design of new communication technology in a safety-critical context, so as to support users even when their behavior is different from that expected, while still ensuring flexible cooperation mechanisms.

When a method addresses a broad range of issues it can require considerable effort to be applied. To decrease this effort some tool support can be helpful. Thus, further work on developing tool support for the proposed method is foreseen. Further enhancements will assist in the identification of possible deviations and problematic traces, and in the discovery of design solutions that cope with them.

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers and Wayne Gray for their helpful comments.

REFERENCES

BAINBRIDGE, L. 1983. Ironies of automation. *Automatica 19*, 775–779.

BENTLEY, R., HUGHES, J. A., RANDALL, D., RODDEN, T., SAWYER, P., SHAPIRO, D., AND SOMMERVILLE, I. 1992. Ethnographically-informed systems design for air traffic control. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work* (CSCW '92, Toronto, Canada, Oct. 31–Nov. 4), M. Mantel and R. Baecker, Eds. ACM Press, New York, NY, 123–129.

BURNS, D. J. AND PITBLADO, R. M. 1993. A modified HAZOP methodology for safety critical system assessment. In *Directions in Safety Critical Systems—Proceedings of the Safety-Critical Systems Symposium*, Springer-Verlag, Vienna, Austria.

CARROLL, J. M. AND ROSSON, M. B. 1992. Getting around the task-artifact cycle: How to make claims and design by scenario. *ACM Trans. Inf. Syst. 10*, 2 (Apr.), 181–212.

CHATTY, S. AND LECOANET, P. 1996. Pen computing for air traffic control. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (CHI '96, Vancouver, B.C., Canada, Apr. 13–18), M. J. Tauber, B. Nardi, and G. C. van der Veer, Eds. ACM Press, New York, NY, 87–94.

FIELDS, R. E., HARRISON, M. D., AND WRIGHT, P. C. 1997. THEA: Human error analysis for requirements definition. Tech. Rep. YCS-97-294. Department of Computer Science, University of York, York, UK. Available via http://www.cs.york.ac.uk/ftp_Hlt468789293d_Hlt468789293ir/reports/.

GRAY, W. D., JOHN, B. E., AND ATWOOD, M. E. 1993. Project Ernestine: Validating a GOMS analysis for predicting and explaining real-world performance. *Human-Comput. Interact. 8*, 3, 237–309.

HOLLNAGEL, E. 1993. *Human Reliability Analysis—Context and Control*. Academic Press, Inc., New York, NY.

HOPKIN, V. D. 1988. Air traffic control. In *Human Factors in Aviation*, E. L. Wiener and D. C. Nagel, Eds. Academic Press, Inc., New York, NY, 639–663.

HUGHES, J. A., RANDALL, D., AND SHAPIRO, D. 1992. Faltering from ethnography to design. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work* (CSCW '92, Toronto, Canada, Oct. 31–Nov. 4), M. Mantel and R. Baecker, Eds. ACM Press, New York, NY, 115–122.

JOHN, B. E. AND KIERAS, D. E. 1996. Using GOMS for user interface design and evaluation: which technique?. *ACM Trans. Comput. Hum. Interact. 3*, 4, 287–319.

KIRWAN, B. 1994. *A Guide to Practical Human Reliability Assessment*. Taylor and Francis, Inc., Bristol, PA.

LEATHLEY, B. A. 1997. HAZOP approach to allocation of function in safety critical systems. In *Proceedings of the 1st International Conference on Allocation of Functions* (ALLFN '97, Galway, Ireland), IEA Press.

MACLEAN, A., YOUNG, R. M., BELLOTTI, V. M. E., AND MORAN, T. P. 1991. Questions, options, and criteria: Elements of design space analysis. *Human-Comput. Interact. 6*, 3-4, 201–250.

MACKAY, W. E., FAYARD, A.-L., FROBERT, L., AND MÉDINI, L. 1998. Reinventing the familiar: Exploring an augmented reality design space for air traffic control. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (CHI '98, Los Angeles, CA, Apr. 18–23), M. E. Atwood, C.-M. Karat, A. Lund, J. Coutaz, and J. Karat, Eds. ACM Press/Addison-Wesley Publ. Co., New York, NY, 558–565.

NIELSEN, J. 1993. *Usability Engineering*. Academic Press Prof., Inc., San Diego, CA.

PATERNÒ, F. 1999. *Model-Based Design and Evaluation of Interactive Applications*. Springer-Verlag, Vienna, Austria.

PATERNÒ, F., SANTORO, C., AND TAHMASSEBI, S. 1998. Formal models for cooperative tasks: Concepts and an application for en-route air traffic controlP. In *Proceedings of the Conference on Design, Specification, and Verification of Interactive Systems* (DSV-IS '98), Springer-Verlag, Vienna, Austria, 71–86.

PATERNÒ, F., SANTORO, C., AND FIELDS, B. 1999. Analysing user deviations in interactive safety-critical applications. In *Proceedings of the Conference on Design, Specification, and Verification of Interactive Systems* (DSV-IS '99), 189–204.

PERROW, C. 1984. *Normal Accidents: Living with High Risk Technologies*. Basic Books, Inc., New York, NY.

REASON, J. 1990. *Human Error*. Cambridge University Press, New York, NY.

SCHMIDT, K. AND SIMONE, C. 1996. Coordination mechanisms: Towards a conceptual foundation of CSCW systems design. *Comput. Supp. Coop. Work 5*, 2-3, 155–200.

SYMON, G., LONG, K., AND ELLIS, J. 1996. The coordination of work activities: Cooperation and conflict in a hospital context. *Comput. Supp. Coop. Work 5*, 1, 1–31.

UK MINISTRY OF DEFENCE. 1996. HAZOP studies on systems containing programmable electronics. Interim Def. Standard 00-58, Issue 1. UK Ministry of Defence.

WHARTON, C., RIEMAN, J., LEWIS, C., AND POLSON, P. 1994. The cognitive walkthrough method: A practitioner's guide. In *Usability Inspection Methods*, J. Nielsen and R. L. Mack, Eds. John Wiley and Sons, Inc., New York, NY, 105–140.

WOODS, D. D., JOHANNESEN, L. J., AND COOK, R. I. 1994. Behind human error. SOAR Rep. 94-01. CSERIAC.

Received: December 1998; revised: July 1999 and October 1999; accepted: November 1999