Is Paper Safer?
The Role of Paper Flight Strips in Air Traffic Control

Wendy E. Mackay

University of Aarhus
Åbogade 34
DK-8200 N, Aarhus, Denmark
mackay@daimi.au.dk

ABSTRACT

Air traffic control is a complex, safety-critical activity, with well-established and successful work practices. Yet many attempts to automate the existing system have failed because controllers remain attached to a key work artifact: the paper flight strip. This article describes a four-month intensive study of a team of Paris en route controllers in order to understand their use of paper flight strips. The article also describes a comparison study of eight different control rooms in France and the Netherlands. Our observations have convinced us that we do not know enough to simply get rid of paper strips, nor can we easily replace the physical interaction between controllers and paper strips.

These observations highlight the benefits of strips, including qualities difficult to quantify and replicate in new computer systems. Current thinking offers two basic alternatives: maintaining the existing strips without computer support and bearing the financial cost of limiting the air traffic, or replacing the strips with automated versions, which offer potential benefits in terms of increased efficiency through automation, but unknown risks through radical change of work practices. We conclude with a suggestion for a third alternative: to maintain the physical strips, but turn them into the interface to the computer. This would allow controllers to build directly upon their existing, safe work practices with paper strips, while offering them a gradual path for incorporating new computer-based functions. Augmented paper flight strips allow us to take advantage of uniquely human skills in the physical world, and leave the user interface and its subsequent evolution in the hands of the people most responsible: the air traffic controllers themselves.

KEYWORDS

Air traffic control, ethnographic study, paper flight strips, safety factors

INTRODUCTION

Air traffic control is a classic example of a safety-critical system involving high risks. Controllers hold the fates of thousands of people in their hands; mistakes that result in crashes are simply not acceptable. Theirs is a complex, collaborative activity, with well-established and successful work practices, requiring rapid responses to constantly changing conditions.

Today's international air traffic control system was introduced shortly after the second World War by the American military and has not changed substantially since the 1960's. Controllers use radio to communicate with pilots and telephones to communicate with each other. The RADAR provides a two-dimensional representation of aircraft moving along pre-defined routes within an air sector, while paper flight strips allow controllers to track and modify information about aircraft and their flight plans (see Hopkin, 1995, for an excellent summary). Of course, the system has also evolved over the years. RADAR and other technologies have been improved and flight sectors, which define the routes through the air space, are reorganized regularly to meet changing traffic patterns. Even so, the basic user interface and corresponding work practices have remained the same, with relatively minor variations at the country, flight center, sector, team and individual level.

Like most western countries, the French government has invested heavily over the past decade to update their current air traffic control system. One of the reasons is technical: much of the existing technology is old and increasingly difficult, and sometimes impossible, to replace. If hardware systems must be replaced anyway, it makes sense to take advantage of the tremendous advances in computing achieved over the past few decades. The second reason is both ethical and political: levels of air traffic are increasing rapidly and the current system is extremely, but not completely, safe. According to a report in the International Herald Tribune (15 January 1997), if current safety levels are maintained and air traffic continues to increase at its present rate, an aircraft will crash somewhere in the world *every week* by the year 2004. Airline management and government officials are understandably concerned that the public will perceive this as unacceptable. They must either increase the level of safety or limit the number of flights. Given the economic consequences of the latter, it is not surprising that numerous attempts are underway to re-examine and improve the current system.

Yet improving air traffic control presents an interesting user interface design challenge because the existing system is already extremely safe. In France, for example, no fatalities have ever been attributed to civilian controllers.[1] New tools must not only offer improvements, but also avoid generating problems. A tool that increases controller efficiency cannot also be permitted to lower safety. A tool that better supports controllers in a crisis cannot reduce their vigilance during slow periods. All new systems must be introduced within the context of the existing environment and help controllers find the optimal balance between aircraft safety and the smooth flow of traffic.

Over the past few decades, controllers have accepted improvements to the RADAR (relatively)[2] easily. They have, however, resisted most attempts to "improve" the other key artifact, paper flight strips. One reason is that nobody has tried to get rid of the RADAR, whereas virtually all new systems replace or create electronic simulations of flight strips. Controllers are probably right to be cautious: the history of automation is filled with examples of expensive new computer systems that either reduced user productivity or were discarded completely (Suchman, 1987, Zuboff, 1988, Dertouzous, 1990,

---

[1] However, a mid-air collision occurred in 1973, when military air traffic controllers took over temporarily during a civilian controller strike. The strike was quickly resolved and French controllers have had significantly more power and better working conditions since then.

[2] In France, in 1965, controllers were given a keyboard to enter modifications in flight plans. The same technology also enabled them to see airplane flight codes on the RADAR (instead of tiny, undifferentiated blips). They demanded and received a bonus for using this new technology. According to one controller, *"If we had not received something interesting in exchange, we never would have done it."* (Poirot-Delpech, 1991).

Landauer, 1995). Unlike office-based users, who often must accept whatever technology is thrust upon them, air traffic controllers have a real voice in the technology they use. Because of the safety-critical nature of the system, they can reject interfaces they do not like. Controllers have the final say for a very simple reason: if there is an accident, computers do not go to jail, controllers do.

Initial perspectives on controllers and paper flight strips

This research project was sponsored by the Centre d'Études de la Navigation Aérienne, the French national center for research on air traffic control. The initial research question was: How can we help air traffic controllers make the transition from today's paper flight strips to "more modern" computer-based systems? We began by talking to people in the organization, to learn how they characterize the work of controllers, in particular the role of paper flight strips in that work. Perspectives varied according to each person's background, although frustration was shared by everyone. Many engineers, both developers and researchers, reported that the controllers did not appreciate the innovations offered by the new hardware and software. Many of the ergonomists and social scientists felt that their findings were not adequately incorporated into the engineers' designs. Controllers who tested new systems often felt that their voices had not been heard and found the new systems to be both slower and potentially more dangerous. When we probed deeper, we found fundamental differences in how each group characterized the day-to-day work of the controllers and, in particular, their use of strips. These characterizations greatly influenced how they designed, evaluated, or reacted to proposed new user interfaces.

Software engineers base their understanding of air traffic control on official rules and written documentation, as required by the International Civil Aviation Organization. They divide air traffic control into sequences of discrete, rational tasks. Controllers are said to have a series of goals and sub-goals, using rules to decide which tasks to perform in order to accomplish each goal. The engineers model this collection of tasks and sub-tasks on the computer to provide the foundation for automation or semi-automation of the controllers' work (e.g. Zorola-Villarreal et al., 1995). Engineers usually separate the tasks and responsibilities for each controller, simplifying the software design but also reducing the cooperative nature of the work. Their designs emphasize coordination across discrete roles, rather than support for cooperative, overlapping activities. Engineers we spoke to felt that paper flight strips are important because of the information printed on them, which was reflected in their software designs. An interesting exception is research by Chatty & Lecoanet (1996), which codified annotations based on an analysis of strips by Preux (1994), to allow controllers to communicate with the computer via written gestures. Engineers want to capture currently inaccessible information, such as flight level changes that are hand-written only on paper strips, in order to create automated systems that increase safety, efficiency or both. Many believe that one of the controller's roles should be to act as a source of input for the computer system, to increase the accuracy of conflict detection and other on-line tools. They are usually surprised when controllers object to data-entry tasks that serve no direct function in controlling the traffic. Many interpreted this as evidence of resistance to change in general and to computers in particular.

One engineering manager remarked that paper flight strips are "old-fashioned"; controllers should be induced to use "modern" systems. (Note that this equates the mouse, keyboard and monitor interface with using computers.) Paper strips are seen as a historical artifact that can be replaced either with an on-line version (Leroux, 1993, Bressolle et al., 1995) or removed entirely (Vortac & Gettys, 1990, Ammerman & Jones, 1988, Bentley et al., 1992). Removing paper flight strips would "relieve controllers of time-consuming activities that do not enhance job performance", according to Edwards et al. (1995). Their new interfaces incorporate only the information contained in strips and they argue that there is no functional difference between writing on a strip and updating an electronic image of a strip by selecting a menu item with a mouse.

Cognitive Ergonomists are another important group at the C.E.N.A. Their role is to take human factors into account when designing new systems. Like the engineers, they describe air traffic control as a series of cognitive tasks, with goals and subgoals (Gras et al., 1994). They build models (Amaldi, 1993, Bressolle et al., 1995) using theories from Cognitive Psychology about decision making (e.g., Kahneman et al., 1982) and analysis of human error (e.g., Reason, 1990). Their research is often conducted in the context of creating requirements specifications for engineers developing new computer systems. Those we spoke to shared the engineer's assumptions about the advantages of "modern" systems and the need to replace paper strips. An important group focuses on controller error and seek to change the existing work practices to increase safety.

Sociologists, a different group of social scientists, are more interested in the social and historical context of the work. Harper et al. (1991) and Hughes et al. (1992) emphasize how the context of the work is essential for understanding both the controller's activities and the role of paper flight strips. For example, Bressolle et al. (1995) has demonstrated that when traffic levels increase, controllers speak to each other less often and write more on the strips. Poirot-Delpech (1995) argues that strips form an essential part of a controller's identity and play a symbolic role as the physical objects representing the otherwise-invisible aircraft in the air. Hopkin (1995) studied how flight strips have evolved over the past half-century, allowing the controllers to flexibly incorporate on-going changes in the air traffic control system. We did not see evidence that these social and historical accounts were considered in the design of the new systems.

Each of these perspectives has limitations. Software engineers are primarily interested in the new benefits offered by the systems they design; it is difficult for them to assess the importance of intangible safety features built in to the existing system. Cognitive ergonomists who seek to find and prevent errors may undervalue successful work practices and underestimate the risks involved in changing them. Sociologists may offer many interesting insights into the context and practice of air traffic control but, with a few exceptions, e.g., Hughes et al., 1992, rarely influence system design. There is, of course, another important perspective: that of the air traffic controllers themselves.

Air traffic controllers like paper flight strips. The interface is familiar, easy-to-use, helps controllers instantly understand the current state of the traffic and lets them communicate without interrupting each other. For example, Figure 1 shows three controllers working independently but monitoring each others' activities. The closest (planning) controller is reaching for a new strip while checking the RADAR and relaying information to a controller at another flight center on the phone. The middle (RADAR) controller is talking to the pilot of another aircraft and comparing two strips, touching them with his fingers. The third standing (planning) controller has just finished a phone call and is updating information about a particular flight via a touch-sensitive screen called a "Digitatron", which lists incoming and outgoing traffic and lets controllers register "official" changes in flight information.

Paper flight strips are flexible, letting controllers easily accommodate the on-going changes in air traffic, administration or regulations. Perhaps more importantly, strips are reliable and, unlike computers, telephones, radio and RADAR, they do not break down. (Of course, the computers that print strips can and do break down, at which point controllers simply hand-write the strips.) Since controllers are responsible for people's lives, they want a system that works even when all other systems fail.

Figure 1: Three controllers working collaboratively at west sector UX in Athis Mons. Behind them, two controllers are working at the adjacent sector.

French controllers are unwilling to act as "input devices" for a computer system of unknown benefit, particularly if data entry tasks distract them or slow them down. Although their over-riding concern is safety, controllers are also responsible for the efficient progress of each aircraft. They are under continuous pressure to both increase the level of air traffic and ensure that no accidents occur. Controllers constantly break the rules in order to accommodate these conflicting goals. One controller claimed that a "work to rule" strike, i.e. exactly following official rules, can reduce traffic by 50% or more. If even partially true, this calls into question the wisdom of using technology to force adherence to abstract rules, rather than incorporating human judgment.

RESEARCH APPROACH

This article describes two field studies undertaken to learn about air traffic control in general and paper flight strips in particular. To gain an in-depth understanding of the use of paper strips, we ran a four-month ethnographic study of a team of controllers at the Athis Mons (Paris) en route control center. To gain a broader understanding of the diversity of the use of paper strips, we conducted a comparison study of eight air traffic control centers, in France and the Netherlands. The next two sections describe these field studies with a summary of the results and a discussion of usability and safety issues. We conclude by suggesting a new strategy for updating air traffic control systems: Augmented flight strips.

FIELD STUDY 1: ETHNOGRAPHIC STUDY OF A TEAM OF AIR TRAFFIC CONTROLLERS

We were invited to spend several months at the Athis Mons *en route* control center, near Paris's Orly airport. We studied team 9-West, following their schedule for four months, including nights and weekends, and observed them under the full range of traffic conditions.

Research Setting: The Athis Mons center is responsible for some of the most complex air traffic in Europe, covering approximately one-fifth of France. Unlike "tower" control centers, which handle take-off and taxiing, en route centers handle high-altitude travel between airports. Civilian and military flights fly overhead in all directions and must be coordinated with aircraft leaving and approaching Paris’s two major airports, Orly and Roissy-Charles de Gaulle. Athis Mons has 21 separately-controllable sectors, half designated as "east", the rest as "west". Controllers are qualified to manage all 10 or 11 sectors in either east or west, but not both. Sectors are combined in different configurations. A single control position can handle a single sector (during peak periods), a group of sectors (during moderate traffic conditions through the day), or all sectors in east or west (throughout the night). The control room was under construction during our study, increasing the number of power failures and overall noise level.

Participants: During our study, team 9-west had four senior controllers, six qualified controllers and five students. French teams are self-managed. Instead of an official supervisor, senior controllers take turns acting as "chef de salle" when the team is on duty. Controllers are responsible for managing the air traffic: they rely upon flight plans, requests by pilots, requests from other sectors, current weather and local traffic conditions to organize the planes and to judge the safest and most efficient ways for aircraft to proceed through the air space. Normally, two or three controllers work together at a control position, managing from 1 to all 11 flight sectors. The RADAR controller manages up to 20 aircraft and negotiates with the pilots. The planning controller receives and organizes incoming traffic and coordinates with controllers at other sectors. As traffic becomes heavier, groups of sectors are split apart and handed off to separate control positions (a degroupment). When traffic becomes lighter, sectors are grouped together into a single control position (a regroupment).

Student controllers usually take three years to become qualified, dividing their time between the classroom and apprenticing in the control room. They spend as many hours as possible controlling the traffic. So much in fact that it is rare to find a position controlled exclusively by senior controllers.

Teams always have visitors, such as instructors being re-qualified. Permanent members of the team leave for six months at a time, for additional training or teaching. Controllers work for two or three days, followed by two or three days off, in an on-going cycle. This ensures that they share the load of working nights, weekends and holidays. (The schedule has now changed slightly, so that controllers always work three days on, three days off.) The work day varies, from as few as five hours to as many as twelve hours, when they work through the night. On any particular work day, controllers go on duty for at most two hours before taking at least a half hour break.

Data Collection: For the first six weeks, we simply observed the controllers at work. We explained that we were not trying to become controllers *per se*, but to understand how controllers use their tools, particularly paper flight strips. We split our time between informal conversations and focused observation. After the first six weeks, we began to systematically videotape 60-90 minute sessions. Table 1 shows the range of sessions, organized by day of the week and time of day. The two letter code indicates the sector observed. (Some sessions have two sectors, indicating a regroupment or degroupment.) We concentrate on the busiest days of the week, Monday, Thursday and Friday, but also included at least one video session of every weekday. We also balanced our coverage over the course of the day, concentrating on the busiest times, but also including slow times. This data covers a wide range of situations, including sunny to stormy weather, with and without strikes, and a range of calm to very stressed conditions.

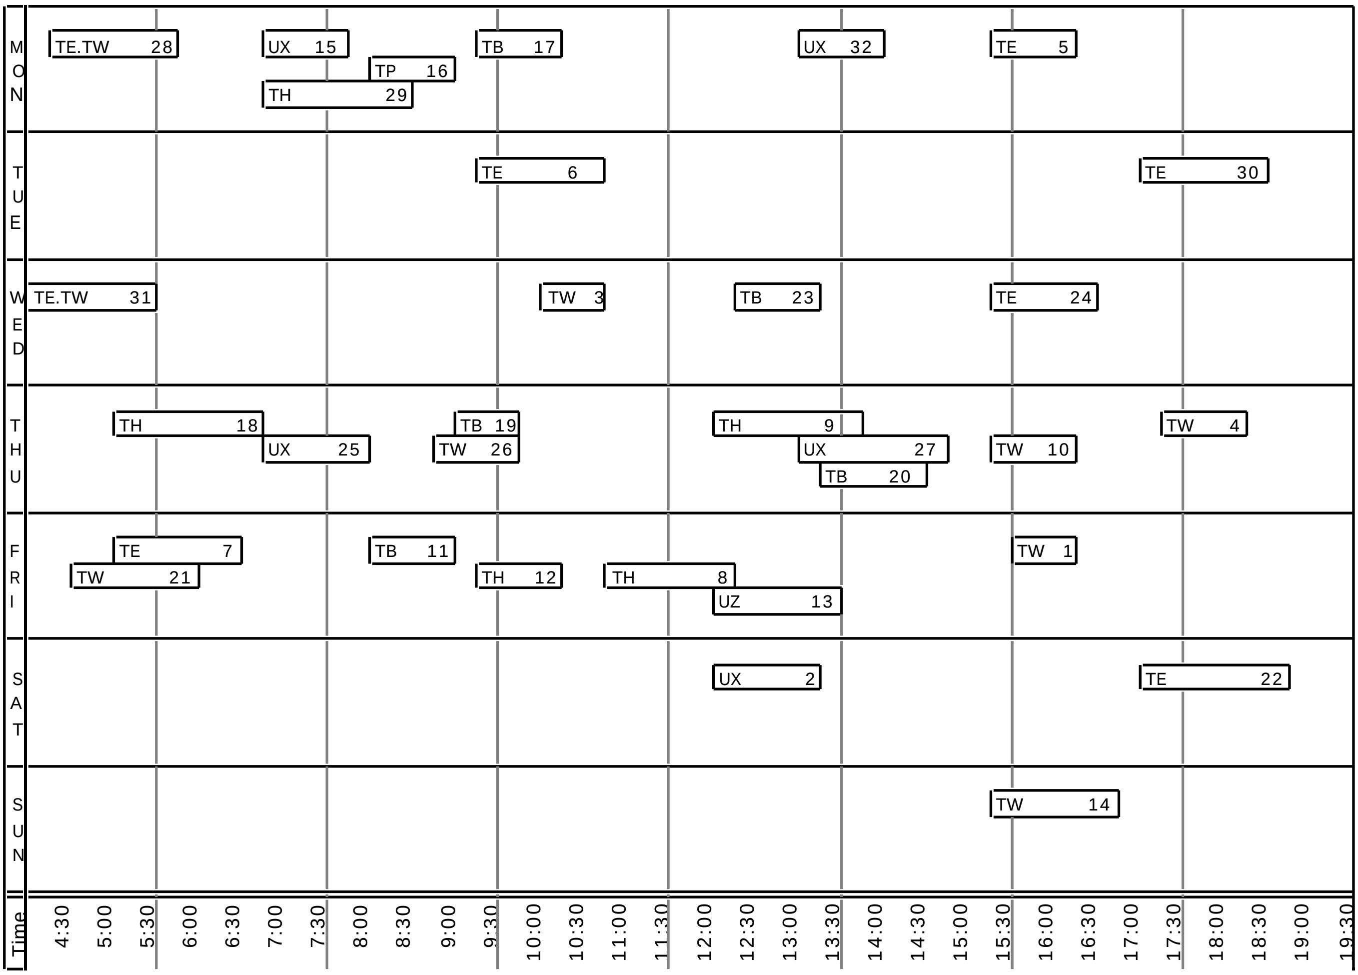

Table 1: Videotape summary by date, time, sector, and length.

We made detailed, timed notes of over 100 hours of work, including almost 50 hours of video. We counted certain activities, paying particular attention to communication patterns and use of paper flight strips, RADAR, Digitatron, radio and telephone. We also noted events that occurred out of the view of the camera, both during taping and at other times. For example, we witnessed several "near misses", in which two aircraft came too close to each other, power failures, and misunderstandings between pilots and controllers. We also periodically recorded the state of all control positions in the room, including number of strips, controllers, and flights listed on the Digitatron screen, to get a picture of the overall traffic patterns and local context during our taping sessions.

We videotaped senior controllers, qualified controllers and students during continuous sessions, regroupments, degroupments and team changes (relief). We were present during training sessions and qualification exams as well as normal operations. Most shots involved a fixed view from a video camera on a tripod, which made it easy to observe and count certain kinds of activities. We also experimented with a hand-held camera, to better capture the overall context of the work. We collected examples of collaborative activities, such as two controllers writing on the strips at the same time and the hand-off among controllers during degroupments, regroupments and team changes.

Our observations and discussions with controllers often raised new questions, causing us to return to the control room after the end of the ethnographic study, to gather more data from the new perspective. Using our written notes and the video as a guide, we selected ten sessions for more in-depth coding and analysis. For three sessions, we copied all annotated strips used during the session and analyzed them with respect to the video and our other notes. Although we performed some quantitative calculations, we were mostly interested in the qualitative aspects of the data.

FIELD STUDY 1: RESULTS

We began our analysis by reviewing videos and notes, and then coding activities of interest. We created summary videos which edited together selected clips of similar activities, such as pointing to the RADAR, rearranging the strips or annotating the strips. Table 2 shows an example of a video summary of situations in which controllers point to strips. The initial columns indicate the tape number, sector, date and time, followed by the duration of the clip. The last column lists the various pointing actions observed. We were particularly interested in cooperative activities, such as when two controllers point to strips at the same time, indicated as "2 hands" in the Notes column. We showed these video summaries to the controllers to get their reactions to our interpretations of their activities.

Tape Identification				Begin	End	Notes: Pointing to strips
07	TE	06-Sep	05:15	05:16:03	05:16:14	point and slide
08	TH	06-Sep	10:59	11:01:15	11:01:20	point and slide
19	TB	26-Sep	09:08	09:10:15	09:10:24	point and slide
22	TE	28-Sep	17:20	17:22:22	17:22:28	point and slide
07	TE	06-Sep	05:15	05:44:16	05:44:26	2 hands
08	TH	06-Sep	10:59	12:04:27	12:04:48	2 hands + point and slide
27	UX	03-Oct	13:19	13:25:53	13:26:09	2 hands
12	TH	13-Sep	09:33	10:26:56	10:27:00	pen sweeps across strips
12	TH	13-Sep	09:33	10:16:24	10:16:32	get rid of 2 strips at same time

Table 2: Video clip summary: Instances of pointing to strips.

The next section begins by explaining two key technologies: RADAR and paper flight strips. We then present our observations about how controllers use strips and other technologies to support their work, and comment on the implications for safe control of the traffic.

Two views of air traffic: RADAR and Paper flight strips

Controllers maintain two complementary views of the air traffic. The first is the RADAR screen, which lets them track the current position of their aircraft. The second is the collection of paper flight strips, which lets them organize the traffic, plan their strategies and record key decisions. Figure 2 shows a RADAR screen under moderate traffic conditions. Each aircraft is represented by a point of light, accompanied by the flight identifier, current speed, flight level and a tail showing recent positions. Some, but not all, routes and beacons are indicated on the background. The RADAR image is two-dimensional, but represents aircraft moving along pre-defined routes within sectors in a three-dimensional space. The controller can modify the image, e.g., zoom in and out or change the level of detail on aircraft labels. Controllers have a large and a small RADAR screen. The large screen is usually set to show the full set of traffic the controller is responsible for; the small screen may be zoomed in or out by different controllers, to examine a particular conflict or traffic situation.

Figure 2: Close-up of the RADAR screen. The motion blur is due to the constant movement of the aircraft on the screen.

The other view, provided by the paper flight strips, is created and maintained by the controllers. Unlike the RADAR view, which changes independently of their actions, the layout of the strips reflects the controller's personal view of the traffic. While this is usually the individual RADAR controller's view, the planning controller sometimes adds, removes or rearranges strips, particularly if senior to the RADAR controller.

Figure 3 shows a typical flight strip from the Athis Mons en route control center. Each flight strip is a record of the flight plan filed by the pilot and authorized by the air traffic control center. Several minutes before each aircraft is due to arrive, a corresponding paper strip is printed with information relevant to that aircraft's planned route through the sector. The strip is divided into discrete sections. The far left section identifies the flight and characteristics of the aircraft, including the flight name and number (Air France flight AFR 540), the RADAR identification number (4332), the type of aircraft (Boeing 737), its speed (450), the departure airport (LFBO = La France, Bordeaux) and the arrival airport (LFPO = La France, Paris' Orly airport). This pilot, when filing the flight plan, requested a flight level of 310. In the next segment, the number 310 is enclosed in angle brackets, indicating that the requested level was approved. The following sector is called Terminal West (TW), indicating who will receive the aircraft next. The next two segments track the altitude (flight level) of the aircraft while they are in the sector. Here, the authorized flight level is 310. The flight will pass by a series of beacons which comprise its route. Here, the aircraft is scheduled to leave the BALAN beacon at 13:13 and pass over AMB, CDN and EPR respectively, at the listed times (13:20, 13:26, 13:31). Note that these are estimates of the planned times, not a record of what actually happened. The section on the far right indicates that this flight strip was printed for sector UX, on 4 October, 1997.

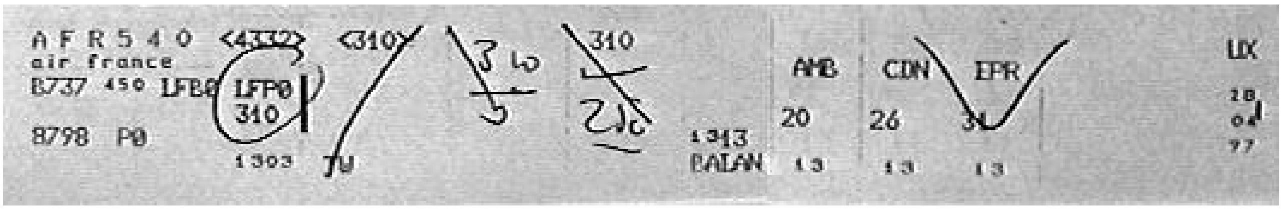

Figure 3: A flight strip describing Air France flight 540.

Controllers annotate strips in predictable ways that are easily interpreted by other controllers. Here, the controller has circled LFPO to show that this aircraft will be handed off to Orly airport approach control, to land. This lets the controller glance along the column of strips and quickly pick out the aircraft that are about to descend towards the airport. The slash mark in the next section (from <310> to TW) indicates that the controller has authorized the approved level and the pilot has agreed. The next two sections show that the flight level has changed from 310 (31,000 feet) to 250 (25,000 feet) while the aircraft has been in the sector. The downward arrow is a reminder that the aircraft is descending. Each time the pilot and controller agree to a new flight level, the controller underlines it. This serves as the legal record of the agreement between the pilot and the controller. The next set of segments indicates the route. Here, the controller has marked a "V" over the EPR beacon, indicating that the pilot is authorized to fly a "direct route" from BALAN to EPR, skipping the intermediate beacons AMB and CDN. This is an example of the kind of safety-efficiency trade-off made by controllers: direct routes save time and are often requested by pilots, but they are inherently more risky. Traffic routes cross each other and aircraft move in both directions; direct routes increase the likelihood of a collision.

Observations about using strips

Many researchers have emphasized the importance of paper flight strips (Harper et al., 1991, Preux, 1994, Hopkin, 1993). Our own observations support their conclusions: paper strips are extremely flexible, take advantage of both visual and tactile memory, and form an essential component of today’s air traffic control system. They also offer many subtle and intangible safety benefits, some of which are articulated below. Undoubtedly, there are other safety benefits (and risks!) we have not seen.

Arguably the most important activity that controllers perform is the continual, sequential checking of each aircraft, first on the RADAR and then on the strips. This routine is important, not only when things are hectic, but also when things are slow. Controllers do this regardless of whether they have a single aircraft or two dozen, whether the traffic is calm or there are multiple conflicts. Such checking is habitual and second-nature to a controller, ensuring that they stay vigilant in low-traffic conditions and that they do not forget other aircraft when solving conflicts in high-stress conditions. Several controllers commented that they were suspicious of computer tools that tried to eliminate this routine during high-stress situations because it would make the work more boring in calm situations and they were afraid that it would be too easy to stop paying attention. Any new tool that fundamentally changes this work practice must demonstrate that the increased safety risk from inattention in low traffic levels is more than offset by increased safety in high-stress conditions.

The physical layout of the strips with respect to the RADAR provides a temporal as well as spatial framework for managing activity. The RADAR provides a global view of the traffic and the strips provide successively more detailed information. Controllers "insert" actions to do in the future, such as remembering to call a pilot in three minutes to authorize a direct route or solving a conflict with a level change, as they work through the cycle. Controllers thus continually move between a global view of the traffic and specific views of individual aircraft or conflict situations. Controllers report that they develop a rich mental image of the traffic during the course of a session. The current strip set up reduces the controller's mental load, allowing them to retain only the important details, since the rest of the information is always instantly accessible in front of them. The physical strips can be viewed as a concrete component of their mental representation, helping them handle more information and successfully deal with interruptions. This has important safety implications for new systems. The current system provides a simple mechanism for controllers to adjust how much or how little of their mental representation is off-loaded onto the strips, through annotations, juxtapositioning of related strips and sliding strips to the side in their holders. If new systems present controllers with too much information per aircraft, reading time and interpretation difficulty will increase, with potential safety implications. If too little information is presented, controllers will be required to increase their mental load to retain the necessary details, which also has potential safety implications. A passive, monitor-based display of strips is unlikely to achieve the right balance and an interactive system under the controller's control would require an extremely light-weight interface to be equally safe.

Physically handling paper flight strips: The notion of affordances (Gibson, 1986) helps clarify some of why controllers interact with strips as they do. Although we usually assume that objects are composed of their physical qualities, Gibson argues that "what we perceive when we look at objects are their *affordances*, not their qualities...what the object affords us is what we normally pay attention to." (p. 134) Paper flight strips are physical objects with multiple affordances that support various aspects of the controllers' work.

When a new strip arrives, the act of removing it from the printer and inserting it into the appropriate strip holder forces the controller to mentally register the new flight. Controllers must physically pick up each strip and place it somewhere; the location determines who will handle it next. Controllers often take strips in their hands as a concrete reminder to deal with that strip next. Controllers periodically reorder the strips during the session. This gives them the sense of "owning" the aircraft and reinforces their memory of the current situation. Layout is important; they place aircraft involved in a conflict next to each other in a column or directly across from each other in two columns. They also use the two columns to separate aircraft into different sectors prior to degrouping. Most controllers, when taking over a control position, physically touch each strip, rearranging some of them. Reordering the strips helps controllers mentally register the new traffic situation. In each case, it is the *act* of rearranging the strips, more than the final layout, that is important. The physical nature of strips also supports cooperative work. The stripboard allows controllers to work independently on different problems within the same collection of strips. The controller who physically picks up a strip or moves it to a new position during a degroupment or regroupment is recognized by everyone as responsible for that aircraft. This has obvious safety implications. The current system permits extensive cooperation and cross-checking among controllers, while maintaining clear individual responsibility. For various reasons, including the mechanics of entering and viewing information, the new system prototypes reduced the overlap in roles among controllers. While this accommodates the need for ascertaining responsibility, it raises the potential of reducing the level of checks and balances in the current system.

Managing strips is a two-handed activity. Guiard's (1987) Kinematic Chain theory has been extremely influential in the analysis of two-handed interaction, particularly skilled bi-manual tasks. Guiard identifies three general principles:

1. *Dominant-to-non-dominant spatial reference:* The non-dominant hand sets the frame of reference for the actions of the dominant hand. For example, for a right-handed person, the left, non-dominant hand is used to position the paper while the right dominant hand actually writes.
2. *Asymmetric scales of motion:* The two hands operate asymmetrically over space and time. For example, the paper-positioning movements of the non-dominant hand are slower and less frequent than the dominant hand as it writes.
3. *Precedence of the non-dominant hand:* The non-dominant hand begins the co-operative bimanual task, follwed by the dominant hand. For example, the non-dominant hand is used to position the paper before the dominant hand begins to write.

We noted that the dominant hand is usually used for writing and precise pointing tasks while the non-dominant hand is used for positioning and framing, as would be predicted by Guiard. Yet, like Balakrishnan et al. (1999), we found some variance from the theory as well. Controllers also shift to use of a single hand, which lets them perform several actions at the same time. A typical sequence involves a controller locating a strip with the non-dominant hand, looking down and beginning to write with the dominant hand, and then looking upward at the RADAR while finishing the annotation and using the other hand to pick up the phone. Controllers usually point to the RADAR with the dominant hand while locating the relevant strip with the non-dominant hand. Controllers sometimes use both hands together, sliding them down both sides of the stripboard as they review the set of flights or look for a particular flight. They usually stop, resting a finger on the relevant strip. Student controllers can be observed "thinking out loud with their hands" as they touch each individual strip involved in a particular conflict. Manipulation of strips varies under changing traffic conditions. Interestingly, controllers are more likely to annotate strips in medium traffic levels than in stressed situations. As traffic grows complex, controllers speak less often and reduce writing to a minimum. They continue to touch strips under all conditions, moving their hands systematically over the stripboard.

Any factor that decreases speed or results in unnecessary shifts of attention is likely to decrease safety. Although two-handed interfaces have been developed in research laboratories, e.g. Bier et al. (1993) and have been reported as much as 40% faster than one-handed interfaces (Buxton and Myers, 1986) they have been mostly ignored in the development of new air traffic systems. Systems that rely upon the mouse/keyboard/monitor found in office-based user interfaces risk being significantly slower than systems that incorporate both hands.

Annotating paper strips: Controllers learn specific rules for how to annotate strips, but individual styles as to how much to write and when vary greatly. Students usually write the most, both to help learn and to demonstrate that they know the rules. A few senior controllers write very little, usually annotating strips quickly just before a new team arrives. However, even controllers with a reputation for not writing must make a minimum number of annotations. Different controllers may indicate the same thing with different annotations. For example, a direct route can be marked with a "V" as in Figure 3, with a line and a loop, as in Figure 4 or with a circle, as in Figure 7. Different teams have different opinions about different types of annotation. For example, one controller was upset that a controller from another team objected to his use of "W" to indicate a warning. In general, one team is in charge and the other teams must follow their rules. One could argue that it would be safer if all controllers annotated strips in exactly the same way, and use technology to enforce the annotation strategy. But individual controllers do differ and systems that treat all controllers as identical may end up creating safety problems, by inadequately accommodating these individual differences.

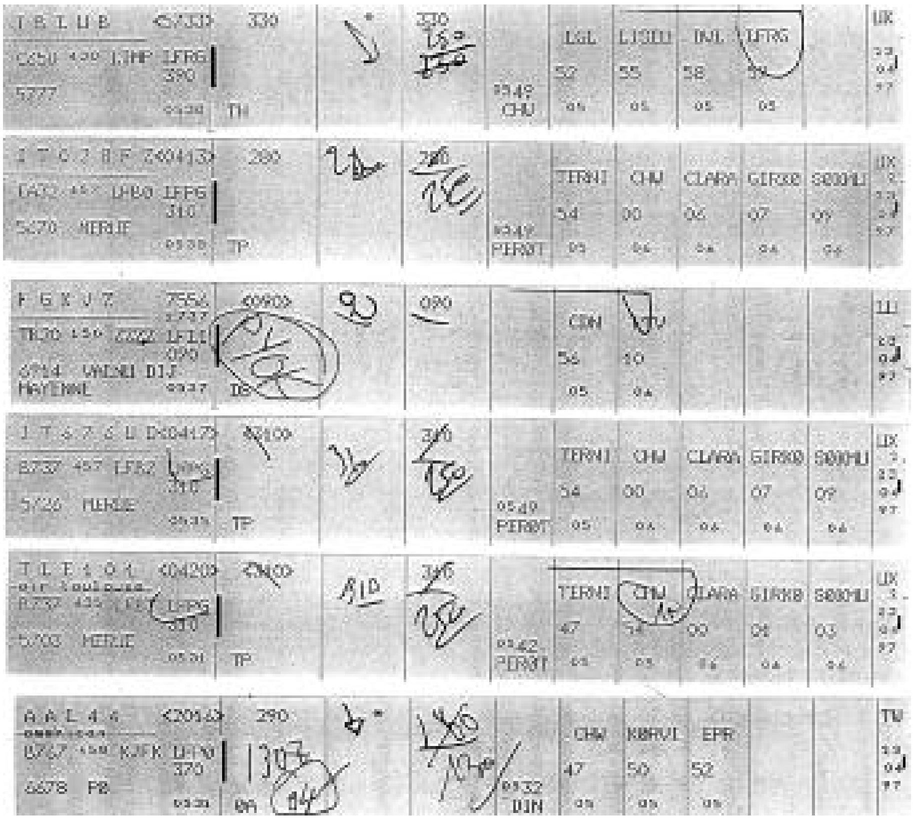

Figure 4: Annotations on a set of flight strips.

Paper strips are very flexible. Controllers hand-write strips in sectors with parachute planes, which do not file regular flight plans and may be in the air for several hours. The controller writes flight information on the back of an ordinary strip, monitoring when it is safe for the parachutist to jump. Hand-written strips are also used to track last-minute changes in military flights or re-route commercial flights. Hand-written strips are also used when printers break down, a relatively common occurrence. Any new computer system must be able to accommodate these and other unplanned control situations, safely and efficiently.

Peripheral awareness

Peripheral awareness, as described by Heath and Luff (1991) in their studies of the London Underground, or situational awareness, as described by Endsley (1988) in her study of pilots, is extremely important for air traffic controllers. They must be able to monitor a variety of auditory, visual and even tactile input , instantly shifting their attention whenever it is required.

Although system designers often try to rid new environments of "extraneous noise", it is important to recognize that sound carries important information. Controllers use the general noise level to get a sense of the overall situation. This is a U-shaped curve: Noise levels are lowest during low traffic periods, rising as traffic increases and controllers chat with each other. Yet peak traffic periods are less noisy: everyone is working and talking is limited to essential conversations. Controllers also learn to track specific noises. For example, at each position, controllers subconsciously prepare for the arrival of a new aircraft when they hear the sound of the strip printer. As in the Heath and Luff study, controllers rely on overhearing each others' phone calls. In fact, planning controllers explicitly talk to two audiences, the controller on the other end of the phone and the adjacent RADAR controller. As mentioned earlier, other "off-duty" team members sit nearby so they can unobtrusively monitor the situation and pitch in as needed. This subtle attention to ambient sound supports both safety and efficiency: only the necessary number of controllers need work on a position at a time, with extra hands instantly available as needed.

Visual cues are also important. The various displays around the control room are designed to give controllers an instant picture of current conditions. For example, the Digitatron has two columns devoted to incoming traffic; when these columns begin to fill up, a new wave of traffic is about to arrive. The display panel above the RADAR has a light for each sector currently controlled by that position; when the light flashes, it indicates not only that a pilot is speaking, but also on which frequency and sector. Controllers glance at the quantity of strips and corresponding level of annotations to get a sense of the traffic. For example, during a storm or when the military have closed a section of the air space, controllers must reroute all aircraft, which results in an easily identifiable pattern of annotations on the strips. These visual cues rely on the human visual system's use of focused and peripheral views. New systems that try to place the same quantity of visual information into a single, focused display, are likely to be more difficult to read and thus less safe.

Tactile cues are subtle. Running one's hands along the strips helps to mentally count them, even when looking elsewhere. We watched student controllers staring at the RADAR, trying and failing to insert a 14th strip into a column that holds only 13. Senior and qualified controllers never make this mistake.

Controllers share a small physical space, which helps them monitor each others' activities. Figure 5 shows two controllers simultaneously annotating two different strips. Each is aware of the others' annotation and will check it at the next opportunity. At night, when all sectors are combined into two control positions, the east and west controllers always sit next to each other. This is an important safety check as they keep each other company and ensure that neither falls asleep.

Figure 5: Two controllers simultaneously annotating two different flight strips (Athis Mons).

Peripheral awareness helps explain a related phenomenon among team members. In light-to-moderate traffic conditions, members of the team who are not assigned to a particular position chat with each other near their team's working control positions. (In contrast, members of the team physically leave the control room during official breaks, after ensuring that the relief team is fully operational.) To external eyes, they appear to be off-duty and perhaps even annoying their "working" colleagues by generating extraneous noise. (A visiting American pilot immediately jumped to this conclusion, for example, and commented on the relaxed life of French controllers.) But this view is mistaken. Although ostensibly ignoring the situation, these controllers are actually tracking what is going on and ready to assist as needed. This helps explain why working controllers rarely need to ask for help; their colleagues usually just pitch in. We discovered that this is a consciously-honed skill. During our interviews in the control room, the controller we were talking to would sometimes stop mid-sentence and turn to pick up the phone or address a problem at a nearby position. After a time, from a minute to half an hour, the controller would turn back and complete the sentence. When we marveled at this, having long since lost track of the thread of the conversation, a senior controller explained that they do this as a form of on-going training. Controllers must operate in a highly interrupt-driven environment while maintaining a model of the evolving state of the traffic. Chatting while continuing to use peripheral cues to accurately monitor the situation keeps them sharp and able to process multiple threads of information. Student controllers are expected to participate and gradually learn this skill.

Controlling peripheral awareness to effect action: Controllers do more than passively accept peripheral information. In some situations, they actively manage it. For example, a planning controller is responsible for making sure that the RADAR controller is aware of incoming traffic problems. She overhears the RADAR controllers' conversations with the pilot as well as his conversations with other controllers, locally or on the phone. She can see when he annotates a strip, points to the RADAR or measures a distance with his fingers. The planning controller is aware of the range of the RADAR controller's visual field. She may keep a new strip in her area, outside of his visual field, if it is not relevant. She moves it into his peripheral view if the strip should be dealt with soon, but not immediately. If the problem is urgent, she will physically move it into his focal view, placing the strip on top of the stripboard or, rarely, actually inserting it. Physically leaning into another controller's "personal space" indicates that a strip must be dealt with immediately, as in Figure 6.

Figure 6: Controllers communicate via physical interaction.

The RADAR controller is aware of this strategy and can actively or passively react to objects moving from the peripheral to the focal view. He can actively seek new strips, even if not placed in his focal area. He can also ignore a strip the other controller deems urgent, maintaining awareness but waiting until he reaches a point at which he can absorb the information. Controllers thus actively or passively push or pull information back and forth between their periphery and focus of attention.

Controllers rely on their ability to extract information from the environment as they need it. They do not create a full mental picture of the traffic before they begin to work: relief RADAR controllers rarely take more than a minute at the new position before starting to direct traffic. They usually survey the strips, gaze at the RADAR, sometimes ask a question or two about an annotation, and then take over, when they are sure they have a safe level of understanding of the situation. The other controller remains in the area for about five minutes, clearly available if a problem arises. The relief controller does not begin the session with the same, rich representation of the traffic as the previous controller, but relies on the strips and RADAR to quickly build such a representation over time.

Peripheral monitoring, both active and passive, provide controllers with efficient methods of assessing the traffic situation and coming to each other's aid. It is difficult to quantify the safety benefits of this type of non-verbal physical communication; however, new computer systems that isolate controllers from each other must somehow accommodate the checks and balances that occur naturally in the paper-based system.

Work practices in a social context: Controllers are very aware of each other's performance. They can tell if someone is tired or less efficient or has become "scared". When someone is sick or back from an absence, they are not expected to act as RADAR controllers; instead they spend as much time as necessary in the supporting role of planning controller, until they are back in shape. Few controllers are over the age of 40; this is a young person's job. We observed one senior controller gradually "retire", moving himself out of the control room into a desk job. The process took several months as he stopped taking control at the working positions and began taking on more and more of administrative functions. The other controllers did not, to our knowledge, discuss this change, but simply allowed it to happen. When we asked, they said he must have been scared by something and decided it was time to move on. This face-saving method of adjusting to each individual's ability level may be unique to French control rooms, which operate with teams whose members take care of and protect each other. More hierarchically-organized control rooms, in which each controller works individually, would presumably involve more a abrupt transition when a controller is no longer able to function at an adequate level. Clearly, any new automated system must similarly take these individual differences in capacity into account to ensure the overall safety of the system.

Controllers are not only aware of each other's failings; they are also admirers of each other's skill. They constantly evaluate each other and seek elegant solutions to problems. Interestingly, "elegance" is defined by the culture of the control room or team. For example, Paris controllers avoid changing flight levels whereas Reims controllers admire them. Yet the rules are not fixed; Paris controllers will change flight levels to avoid generating additional or messier conflicts. This notion of elegant solutions may have safety implications, particularly if it makes controllers' behavior more predictable by one another.

Controllers must continually update their knowledge, since the underlying "theory of the sectors" is constantly changing. They often refer to files containing maps, descriptions and procedures, particularly after they have been out of the control room for over a week. In such situations, other controllers do not expect them to be up-to-date and allow them several days to regain their competence before they take on stressful situations. We watched a senior controller, back in the control room after several years of teaching, struggle with relatively simple traffic on his first day in the control room. The other controllers treated him as a senior controller who needed to relearn the current details, and allowed him more personal latitude than they would have given a student in the same situation.

A few senior controllers test their own limits, creating more complex traffic situations than normal, usually with better throughput, in order to maintain their ability to handle unforeseen emergencies. This has safety implications. At any particular point in time, they may be increasing the risk associated with a particular traffic configuration. But in the long run, others rely on their increased ability to handle complex or unforeseen situations that occur naturally.

No matter what level of expertise, a controller is always expected to "do something". Even very junior controllers are expected to pick up the phone, if only to hand it to someone else. We saw a senior controller chastise a junior controller, not because he had to take over when she was unable to handle five aircraft, but because she did not then keep busy by contributing to the shared flow of work. This also has safety implications: controllers control traffic by controlling traffic. The current system allows controllers to act in familiar ways under all circumstances. Proposed systems that require different work practices under low and high traffic conditions will require controllers to master two work styles instead of one and risk problems when controllers must shift between the two states..

A final observation about the senior and qualified controllers: all are extremely conscious of and articulate about their work practices. (This is in sharp contrast to our studies of people in several other professions.) This is, at first glance, surprising, since the *activity* of controlling aircraft limits talking and emphasizes more subtle non-verbal cues and peripheral awareness. It is their on-going role as teachers, constantly evaluating and explaining the practice of controlling traffic to apprentice controllers, that keeps them aware of their work. After each session, or during slow periods, senior controllers debrief students, quizzing them or providing explanations matched to their particular level of skill. (This may explain their willingness to patiently explain and re-explain situations to us: we were simply treated as very junior controllers.) Their thoughtful examination and interpretation of their own work practices, whether about paper flight strips or complex cognitive skills, is an essential factor in the safety of the current system. New computer systems that automate work and thus reduce the need for controllers to interpret and explain their work practices may pose unknown and difficult-to-measure risks that must somehow be accounted for.

FIELD STUDY 2: COMPARISON STUDY OF EIGHT CONTROL ROOMS

The study at Athis Mons gave us a great deal of information about how strips are managed within a particular air traffic control center. We were also interested in finding out which aspects were common and which were unique to Athis Mons. To better understand the similarities and differences across control centers, we also observed seven other control rooms in France and the Netherlands.

Research Setting: We visited each of three different types of control rooms: en route, tower and approach in both countries. Since our primary concern was en route air traffic control, we decided to compare large and small centers in each country (Paris and Bordeaux in France and Amsterdam and Maastricht in the Netherlands). The tower control centers (at Paris' Orly airport and Amsterdam's Schiphol airport) have a view of the airport and handle take-offs and travel on the ground. The approach centers are also located at the airports, but look more like en route control centers, relying on RADAR rather than windows to guide aircraft as they land.

Participants: At each control center, a host greeted us, explained the basic operation of the control room and introduced us to other controllers. We spent from several hours and two days at each center, interviewing and observing controllers at work. Visits to Maastricht, Bordeaux and Amsterdam involved both high- and low-stress periods; the visit to Orly was during a low-stress period.

Data Collection: We videotaped 1-2 hours at each center, with corresponding detailed notes and observations. The controllers at each center gave us sample strips, including many with annotations, which enabled us to compare the use of paper strips across centers (Figure 7).

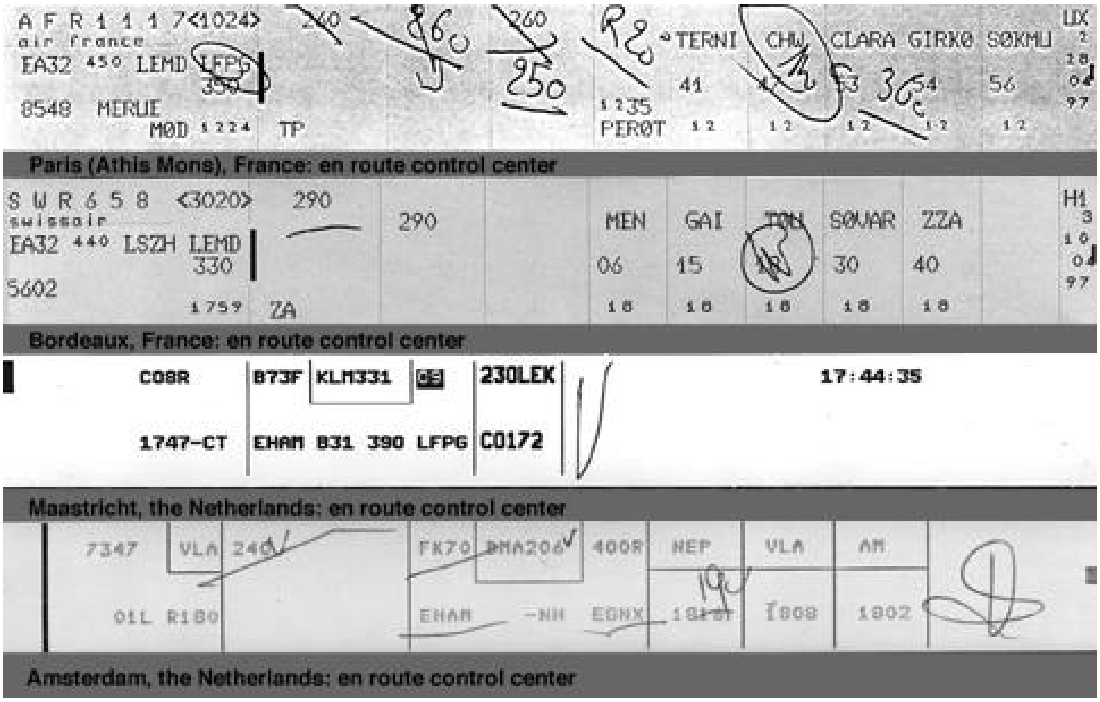

Figure 7: Annotated flight strips from four en route control centers in France and the Netherlands.

FIELD STUDY 2: RESULTS

We examined differences and similarities across the different control centers, paying particular attention to the physical strips. We also looked briefly at cultural differences, not only between countries, but also between large and small centers, and among the three different types of control centers.

Physical strips: Our main concern was the strips themselves, which varied in color (white, yellow, green), size (Maastricht had "mini-strips") and in certain details, e.g., tower centers show a series of flight levels to indicate take-off altitudes. The arrangement of information was clearly customized to meet local needs, yet the basic layout was the same for all. Any controller (and even we!) could immediately recognize and interpret most of what was on each strip. Some annotations were familiar across flight centers, in particular, underlining or indicating a new a flight level. The Dutch centers, which were more hierarchical than the French centers, also reserved space for supervisor's signatures (see the far right initials on the fourth strip in Figure 7).

A major difference across control centers is whether or not an aircraft is represented by a single strip. Tower control centers physically pass, or in some cases, throw, strips from one control position to the next. Each strip's tour around the control tower reflects the aircraft's progress as it taxis through the airport and takes off. In contrast, en route and approach centers print a new strip just prior to the aircraft's arrival into one of the sectors controlled by a particular working position. The only exception is when a control position degroups into separate control positions or regroups by collecting strips from other working positions. In these cases, controllers physically pick up strips from the original position and carry them to the new position. As mentioned before, the physical passing of strips provides a concrete model for allocating and communicating responsibility which is an important safety consideration. The physical nature of strips though, means that they can get lost or mislaid. In fact, we heard two anecdotes about such incidents. In one case, a student controller stuffed a strip in his pocket when he was too busy to integrate it and then spend several uncomfortable seconds wondering where it was when he was ready to insert it into the stripboard. In another case, a tower controller mentioned a case in which one of the thrown strips was dropped and not immediately noticed by the receiving controller. Such situations are apparently extremely rare and the simultaneous presence of the aircraft on the RADAR ensures that "lost" strips are not forgotten. Even so, this is a factor that must be considered in any physical incarnation of strips as the interface to a computer system.

Use of strips: We noted several basic differences across control centers. The most important difference was at Maastricht, which we visited because they were said to have gotten rid of paper strips. We discovered that strips were indeed omitted for some sectors, but only in the simplest traffic sectors. The new system involves an enhanced RADAR screen that includes some information normally printed on strips, as well as a monitor with a system-controlled view of information normally printed on the strips and a touch-sensitive device for inputting information. Although our host, who was not a controller, claimed that the controllers had shifted from paper to electronic strips, we observed that they completely ignored the on-screen image of strips. The only exception in two days was a student controller undergoing his qualification test. When asked, the controllers explained that the electronic strips arranged themselves automatically, “making them useless”.

Maastricht controllers face a very different traffic situation compared to Paris controllers. Each working position controls at most three active sectors, with one-way traffic and a single cross-section in each. In contrast, Paris (west) controllers manage up to 11 active sectors, with bi-directional traffic along the routes and as many as a dozen cross-points in each sector. Aircraft in Maastricht stay in the sector for 2-5 minutes and 80% of the controller’s conversations involved greeting or handing the aircraft off to another sector. Long-term planning is mostly unnecessary as controllers can react directly to the aircraft they see on the RADAR. Paris traffic is very different, with aircraft staying in a sector for an average of 25 minutes or more. Also, the presence of two major airports makes it necessary to handle both aircraft flying over Paris and aircraft going to or leaving two major international airports. Interestingly, despite the relative simplicity of Maastricht, whose controllers were surprised at the complexity of the situation in Paris, we had the sense that those controllers were working harder. Two people were always on duty, actively working, even when traffic was slow. Also, every Maastricht controller had a pad of paper (actually, the back of the old flight strips) and all made hand-written notes. One controller explained that she did not always re-read what she wrote, but the physical act of writing helped her to remember what she had decided. As in Maastricht, our host at the Amsterdam approach center explained how they were trying to get rid of paper strips. He explained that controllers "never write on strips anymore". He cheerfully let us take a set of about 50 used strips, and was surprised when we calculated that 80% of them were annotated. The Amsterdam en route center continues to rely on paper strips, but the approach center, with simpler traffic, has introduced new technology to try to reduce use of paper strips. The result is a proliferation of paper, with hand-written notes on the desk and post-it notes attached to the RADAR.

Controllers in all centers wrote notes to themselves and each other, which raises a safety issue with the new, "paper-less" systems. Is it better to write information in a standardized, clearly-understood way on paper flight strips or encourage an ad hoc, non-standard, less-easily interpretable system that controllers invent to handle unforeseen events?

Layout of strips: Most en route control centers use strip holders (metal or plastic) that fit into two or more metal columns fixed to a desk. Strips can be slid up and down and arranged in separate groups within the columns. Paper strips can also be offset left or right within the strip holder (Figure 8) to "highlight" particular aircraft and to provide an abstract model of the traffic. Offsetting strips provides controllers with a very fast method of indicating conflicts or setting reminders; automated versions that require more effort may be less safe as well as less efficient.

Figure 8: Strips can be offset in both directions.

The Bordeaux en route control center omitted strip holders over a decade ago. They place strips on a table with a series of small steps (Figure 9), which lets them organize aircraft more geographically. The planning controller progressively slides strips leftward toward the RADAR controller, as the aircraft move through the air space. We spoke to two controllers who had worked in both control centers, who explained the trade-offs between the two approaches. The geographical system (without strip holders) works well for Bordeaux, which has lighter traffic conditions but more complex geographical problems: the commercial air space is shared with the French military. The latter can see all the civilian aircraft, but the civilian controllers can see only some of the military aircraft and must defer to the military controllers. Military routes are sometimes made available to civilian traffic but with the understanding that the civilian traffic may be asked to leave with as little as ten minutes notice.

Bordeaux controllers like the visual overview the stripboard table provides. They talk of the board "getting yellow", i.e. filling up with yellow strips, which gives them an instant intuitive feel for the level of traffic. The more abstract system of placing strip holders in columns makes more sense for the

Paris and other en route control centers, since it encourages controllers to group aircraft heading for a particular airport or to organize them based on potential conflicts.

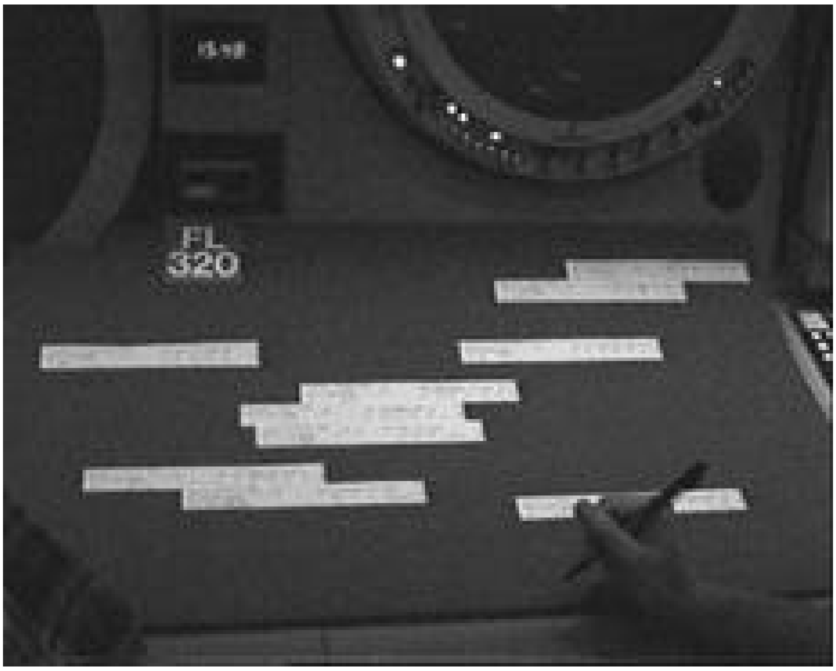

Figure 9: Bordeaux controllers place strips without holders on a table between the RADAR and planning controllers. Strip layout reflects each aircraft's geographical position in the sky.

Differences across control centers: In both countries, the large control centers, which handle more complex traffic, appeared less formal than their smaller counterparts. The smaller centers were more overtly hierarchical, with less informal interaction between controllers and their supervisors. This surprised us: we expected that the greater traffic demands would require more formal organization, but saw the opposite. The French centers were more informal than the Dutch centers, and the most formal center was the smallest, run by EuroControl. This is probably partly cultural and partly due to the organization of the control centers: French controllers work in teams, with shared management by senior team members. The Dutch controllers work as individuals, reporting to a supervisor and working with different controllers at different times. The EuroControl center employs controllers from all across Europe, with many different native languages. English is the common language, but many controllers spoke three or more and would shift languages when talking to different colleagues. This center also had by far the greatest turn-over in personnel and were also the only ones who acted obviously differently when a supervisor was present. They also reported having trouble training sufficient numbers of new controllers, which was not a problem in the other centers.

Tower control centers are physically different from approach and en route centers, with breathtaking views of the airport and beyond. On a clear day, tower controllers have the choice of watching the actual aircraft or the RADAR. The approach center in Amsterdam is proud of being the only center in Europe with windows, which are there for aesthetic rather than work-related reasons. The particular details of the layout and use of paper strips reflects a combination of history, air space and the culture of each control center. Controllers cannot change the RADAR, nor can they change the software tools given to them. But they can and do change the details of their interactions with strips, which helps them adjust to on-going, externally-imposed changes.

IMPLICATIONS FOR SAFETY-CRITICAL SYSTEMS

The design of safety-critical systems differs from that of other interactive systems: While improving productivity is important, safety remains the over-riding concern. Increasing the former at the expense of the latter is simply not acceptable. But what are the factors necessary for a safe system? This paper explores the role of a single physical artifact, the paper flight strip, with respect to the practice of air traffic control. Our findings support and extend the findings of other researchers who have studied flight strips. The current paper-based system supports safe and effective work practices and offers a level of flexibility difficult to imagine with traditional computer-based interfaces. However, the factors that contribute to safety are largely unmeasured and unnoticed and are not part of the official guidelines used to design new air traffic control systems. The dismal record of so many automated air traffic control systems may be attributed, at least in part, to their inability to support these intangible factors.

The title of this paper asks "Is paper is safer?". There is, of course, no single answer. But it does raise the question as to whether the current strategy of radically changing current work practices and eliminating paper strips may be misdirected. The steady rise in air traffic is the reason most often given for automating the air traffic control system (although a level of discomfort with the old-fashioned nature of paper is also apparent). The alternative, to continue using paper strips and artificially restrict the growth of air traffic, is generally viewed as politically and economically untenable. Yet despite vast expenditures of time and resources, automated and partially-automated air traffic control systems have failed to live up to their promise.

This paper suggests a third alternative, which is influenced by two key observations. First, today's air traffic control system is extremely safe and efficient. Attempts to radically change work practices that have successfully evolved over the past 50 years will almost certainly fail to account for all the embedded, intangible safety factors and are likely to result in dangerous, perhaps fatal, situations. Second, adding computer support is not equivalent to adopting the input and output devices currently associated with office automation systems, i.e., a keyboard, mouse and monitor. Alternative input and output devices, including paper itself (Negroponte, 1997, Gibbs, 1998), can provide user interfaces to computer systems that are both familiar to controllers and still provide the benefits of a computer. This approach, called Augmented Reality (Wellner et al., 1993) is still new, but worth exploring. Paper need not be an old-fashioned technology to be tossed away, but rather a new form of computer interface that is truly under the control of its users (Mackay, 1998). Mackay et al. (1998) describe Caméléon, a follow-on project to this study, which is designed to preserve as many of the intangible safety benefits of the existing system as possible, while adding links to the computer. The goal is to build upon existing, successful work practices and gradually incorporate additional functionality, giving controllers an active say in the details and the evolution of the interface.

Augmenting paper strips, rather than replacing them, is clearly controversial: It offers a radically different view of how to automate air traffic control systems. However, based on the multitude of intangible safety factors buried in paper flight strips, and the questionable record of other automation approaches, the time has come for a new approach.

ACKNOWLEDGMENTS

A special thanks to the members of équipe 9-West at the Athis Mons en route control center, for their warm welcome and participation in all aspects of our work. In addition to spending an inordinate amount of time answering our questions and explaining how and why they interact with strips, many chose to give up their personal time to participate in workshops and to contribute ideas, critiques and suggestions during our numerous prototyping sessions. We would also like to thank the many other controllers who allowed us to videotape them and patiently answer our questions, at the en route control centers at Athis Mons (Paris), Bordeaux, Maastricht and Amsterdam, as well as the approach and tower control centers at Orly and Schiphol airports.

Anne-Laure Fayard was instrumental in conducting the field studies, particularly in helping translate when my French was inadequate. She helped log the data and engaged in many discussions of the implications of this research. Thanks also to Lionnel Médini, who also accompanied us into the control room and played an essential role in the development of the later Caméléon project. Special thanks to Stéphane Chatty for bringing me to the C.E.N.A. and for his on-going support. Thanks also to Christophe Mertz and the members of the C.E.N.A. research teams in Toulouse and Paris for their participation in workshops and openness to a different approach. Thanks also to Michel Beaudouin-Lafon, Nicolas Roussel, Stéphane Conversy, and Paul Dourish for discussions of earlier versions of this paper.

REFERENCES

Amaldi, P. (1993) RADAR controller's problem-solving and decision-making skills. In *Verification and Validation of Complex Systems: Additional Human Factors Issues*. Wise et al. (Eds.) Embry-Riddle Aeronautical University.

Balakrishnan, R. Fitzmaurice, G. Kurtenbach, G., & Buxton, W. (1999) Digital Tape Drawing. In Proceedings of the 12th Annual Symposium on User Interface Software and Technology, CHI Letters, Vol. 1, Issue 1, pp. 161-169.

Bentley, R., Hughes, J.A., Randall, D., Rodden, T., Sawyer, P., Shapiro, D. & Somerville, I. (1992) Ethnographically-informed systems design for air traffic control. In *Proceedings of CSCW '92, ACM Conference on Computer-Supported Cooperative Work.* (pp. 122-129) Toronto, Ontario: ACM Press.

Bier, E., Stone, M., Pier, K., Buxton, W. & DeRose, T. (1993) Toolglass and Magic lenses: The see-through interface. In *Procceedings of ACM SIGGRAPH'93*. New York: ACM Press, pp. 73-80.

Bressolle, M.C., Pavard, B., & Leroux, M. (1995) The role of multimodal communication in cooperation and intention recognition: The case of Air Traffic Control. In *CMC'95, The International Conference on Cooperative and Multimodal Communication: Theory and Applications.* Eindhoven, The Netherlands.

Buxton, W. & Myers, B. (1986) A study in two-handed input. In *Proceedings of CHI'86 ACM Conference on Human Hactors in Computing Systems*. Boston, MA: ACM Press. pp. 321-326.

Chatty, S. & Lecoanet, P. (1996) Pen Computing and Air Traffic Control. In *Proceedings of CHI'96 ACM Conference on Human Factors in Computing*, (pp. 87-94) Vancouver, British Columbia: ACM Press.

Dertouzous, M. (1990) *Computers and Productivity*. Cambridge, MA: MIT Press.

Endsley, M.R. (1988) Design and Evaluation for Situation Awareness Enhancement. In *Proceedings of the Human Factors Society 32nd Annual Meeting*, 32(1). Anaheim, CA: The Human Factors Society. pp. 97-101.

Gibbs, W. (1998) The Re-Invention of Paper. *Scientific American* (Sept. 1998).

Gibson, J. (1986) *The Ecological Approach to Visual Perception.* Hillsdale, N.J.: Erlbaum Associates.

Gras, A., Moricot, Poirot-Delpech, S.L., and Scardigli, V. (1994) *Faced with automation: The pilot, the controller and the engineer.* Publications of the Sorbonne: Paris.

Guiard, Y. (1987) Asymmetric division of labor in human skilled bimanual action: The kinematic chain as a model. *Journal of Motor Behavior*, 19(4) pp. 486-517.

Harper, R., Hughes, J., & Shapiro, D. (1991) Harmonious working and CSCW: Computer Technology and Air Traffic Control. In *Studies in CSCW: Theory, Practice and Design.* Bowers, J. & Bedford, S., Eds. North Holland: Amsterdam. pp. 225-235.

Heath, C. & Luff, P. (1991) Collaborative Activity and Technological Design: Task Coordination in the London Underground Control Rooms. In *Proceedings of ECSCW'91, The European Conference on Computer-Supported Cooperative Work.* Kluwer Press.

Hopkin, V.D. (1995) *Human Factors in Air Traffic Control*. London: Taylor & Francis.

Hopkin, V.D. (1993) Human factors implications of air traffic control automation. In *Proceedings of the 5th International Conference on HCI.* Orlando, FL. pp. 145-150.

Hughes, J.A., Randall, D. & Shapiro, D. (October, 1992) Faltering from Ethnography to Design. In *Proceedings of CSCW '92, ACM Conference on Computer-Supported Cooperative Work.* (pp. 115-122) Toronto, Ontario: ACM Press.

Kahneman, D., Slovic, P. and Tversky, A. (1982) *Judgment under uncertainty: Heuristics and biases.* Cambridge: Cambridge University Press.

Leroux, M. (1993) The role of expert systems in future cooperative tools for air traffic controllers. In *Proceedings of 7th International Symposium on Aviation Psychology.* (pp.26-29) Columbus, OH.

Landauer, T. (1995) *The Trouble with Computers*. Cambridge, MA: The MIT Press.

Mackay, W.E. (1998) Augmented Reality: linking real and virtual worlds. *Proceedings of ACM AVI '98, Conference on Advanced Visual Interfaces*. L'Aquila, Italy: ACM.

Mackay, W.E., A.L., Frobert, L., and Médini, L. (1998) Reinventing the Familiar: Exploring an Augmented Reality Design Space for Air Traffic Control. *Proceedings of ACM CHI '98 Human Factors in Computing Systems.* Los Angeles, California: ACM/SIGCHI.

Negroponte, N. (1997) Surfaces and Displays. *Wired*, January issue, pp. 212.

Poirot-Delpech, S. (1995) Biographie du CAUTRA. Naissance et développement d'un système d'informations pour la circulation aérienne. Thèse de doctorat de sociologie. Université de Paris I.

Preux, F. (1994) Rôle des strips dans l'activité des contrôleurs. *Sélection Professionnelle* IEEAC. CENA.

Reason, J. (1990) *Human Error* Cambridge: Cambridge University Press.

Suchman, L. (1987). *Plans and Situated Actions*. Cambridge, England: Cambridge University Press.

Vortac, O. & Gettys, C. (1990) Cognitive factors in the use of flight progress strips: Implications for automation. Norman: Univ. of Oklahoma, Cognitive Processes Laboratory.

Wellner, P., Mackay, W. & Gold, R. (1993) Computer-Augmented Environments: Back to the Real World. Special issue of *Communications of the ACM*, 36 (7).

Zorola-Villarreal, R., Pavard, B., and Bastide, R. (1995) SIM-COOP: A tool to analyse and predict cooperation in complex environments. A case study: The introduction of a datalink between controllers and pilots. In *Proceedings of the 5th International Conference on Human-Machine Interaction and Artificial Intelligence in Aerospace*, IHM-AI-AS 95, Toulouse, France.

Zuboff, S. (1988). *In the Age of the Smart Machine.* New York: Basic Books.