

When Paper Meets Multi-Touch: A Study of Multi-Modal Interactions in Air Traffic Control

Cheryl Savery^{1,3}, Christophe Hurter^{3,4}, Rémi Lesbordes²,
Maxime Cordeil^{2,3,4} and T.C. Nicholas Graham¹

¹ School of Computing Queen's University Kingston, Canada K7L 3N6	² DGAC DSNA DTI R&D 7, Avenue Edouard Belin 31055, Toulouse, France	³ ENAC 7, Avenue Edouard Belin 31055, Toulouse, France	⁴ IRIT, Université de Toulouse, 118 Rte de Narbonne Toulouse, France
---	---	--	--

{cheryl.savery, nicholas.graham}@queensu.ca
{christophe.hurter, remi.lesbordes, maxime.cordeil}@enac.fr

Abstract. When multiple modes of interaction are available, it is not obvious whether combining these technologies necessarily leads to a better user experience. It can be difficult to determine which modes are most appropriate for each interaction. However, complex activities such as air traffic control require multiple interaction techniques and modalities. As a result, in this paper, we study the technical challenges of adding finger detection to an augmented flight strip board used by air traffic controllers. We use our augmented strip board to evaluate interactions based on touch, digital pen and physical paper objects. From our user study, we find that users are able to quickly adapt to an interface that offers such a wide range of modalities. The availability of different modalities did not overburden the users and they did not find it difficult to determine the appropriate modality to use for each interaction.

Keywords: Paper computing; augmented paper; digital pen; interactive paper; tangible interfaces; air traffic control.

1 Introduction

In modern air traffic control centres, paper flight strips continue to be used by air traffic controllers because they provide a tangible interface that aids in the visualization of the aircraft under their control. Flight strips allow easy collaboration and sharing of duties and provide an efficient means of recording communication between pilots and air traffic controllers [16]. However, as the flight strips are mere paper artifacts, the information they contain cannot be leveraged to assist air traffic controllers perform more efficiently, and more importantly, more safely. For example, with a paper based system, it is impossible to automatically alert an air traffic controller if a plane is given clearance to descend to a level before that level has been vacated or to warn the air traffic controller if an aircraft overshoots its assigned flight level. Although computer based systems exist, controllers have been reluctant to adopt them. This reluctance is in part due to the superior interaction qualities offered by paper.

New technologies including multi-touch surfaces, Anoto pens [2] and augmented reality (AR) have led to new interaction techniques that have the potential to bridge

the gap between the digital world and paper flight strips, allowing information entered onto flight strips to be recorded digitally, while still allowing air traffic controllers to manipulate the strips in a physical manner. In addition to increasing safety, a digital system can assist the controller, making it easier to locate flight strips and schedule aircraft. For example, when the controller points to an aircraft on the radar screen the associated flight strip could be highlighted, and similarly if the controller points to a flight strip, the plane could be highlighted on the radar screen. This has led us to the development of an augmented flight control system where controllers can fluidly switch between pen and touch interaction on both physical and digital strips. This supports the familiarity and rapid interaction of traditional techniques while providing the safety and collaboration opportunities of digital interaction.

Combining multiple technologies in this way has the potential, however, of burdening users with too much choice. In a given situation, users must determine whether it is better to use a pen on a paper strip, pen on a digital strip, or use touch-based manipulation. It is not obvious that this proliferation of input possibilities necessarily leads to a better user experience. We therefore ask the question of whether, in the domain of air traffic control, more options and modalities are helpful, or merely lead to confusion. We are interested in how users deal with the availability of different input techniques: do they take advantage of the numerous options, or do they settle on a smaller subset of the options? And if they use only a subset, do all users choose the same subset, or different ones?

In this paper, we present an interface that combines multi-touch, Anoto pen, tangible objects and augmented reality and investigate the interactions involved in several tasks typically performed by air traffic controllers. For each task, we developed a variety of possible interaction techniques combining multi-touch, Anoto pen and/or physical manipulation of paper strips. The interactions were initially developed based on sessions observing the existing processes used by air traffic controllers at the Toulouse Blagnac Airport and at an air traffic control simulation centre in Toulouse. The interactions were then evaluated by air traffic controllers and modifications were made based on their feedback. The air traffic controllers then completed a scenario based on the task of flight stack management during which we observed the techniques they selected to use. Semi-structured interviews were used to gain further insight into the controllers' actions and their reasons for preferring different types of interactions.

This paper makes two main contributions. First, our technical contribution addresses the complex challenges of combining of paper strips, digital pen and finger tracking. Our implementation is simple, flexible and allows finger, pen orientation and tangible object detection. Second, we have investigated the impact of combining multiple input techniques. We find that users are able to quickly adapt to an interface that offers such a wide range of modalities. The availability of different modalities did not overburden the users and they did not find it difficult to determine the appropriate modality for each interaction.

The paper is organized as follows. We first present related work on technologies and interactions for combining multiple types of input, specifically focusing on those that involve pen and/or paper. Next, we describe air traffic control centres, followed

by a description our system that combines multi-touch and pen technology with tangible paper objects and augmented reality. We then present a user study in which air traffic controllers tested and evaluated various interaction techniques and discuss the results from these sessions. Next, we discuss how combining different modalities impacts the user experience. Finally, we conclude by describing how this work will be used to guide the development of the next iteration of our augmented strip board.

2 Related Work

Related work comes from two areas. First we look at research into technologies that have been used to create new user interfaces that combine other technologies with pen and paper computing. Following that we look at research involving the types of interactions that are used with these interfaces.

2.1 Paper and Pen Computing

Technologies for interacting with pen and paper can be divided into four broad categories [22]:

- Scanning and interpreting the content of paper documents;
- Identifying and tracking the location of paper artifacts;
- Capturing input; and
- Outputting information onto the paper

Many pen and paper interfaces have been developed that focus solely on capturing input from a digital pen. These applications cover a diverse range of areas such as documenting scientific research [17], filling out paper forms, composing music [25], mapping [6] and medical records [6]. These applications either record information digitally for later transfer to computer or they may display the data as it is entered allowing verification and correction in real-time.

Other applications focus on tracking paper artifacts and using gestures to change the information displayed. Paper Windows [11] uses markers to track the locations several sheets of paper and a series of gestures to manipulate the information displayed. Mouse-light [21] uses Anoto technology to track the position of a small projector on engineering drawing and display augmented information. Similarly, in [19], a handheld display such as a smart phone is used to augment information displayed on paper maps. Markers such as ARTags [3] allow the position and orientation of objects, including paper, to be tracked in three dimensional space. Do-Lenh et al [7] use a combination of ARTags and touch detection to create multi-finger interactions with paper.

Recent research has investigated techniques for combining pen input with multi-touch surfaces by applying transparent Anoto film to an LCD display [10] or a back projected FTIR touch display [15]. Aitenbichler and Schnelle-Walka [1] describe an architecture for combining Anoto, touch and AR based markers. However in order to

detect the markers through the back-projection foil the authors found that they needed to be constructed using reflective foil for the white portions of the markers.

Despite the wide range of applications and technologies that exist for interacting with pens and/or paper, to the best of our knowledge, no applications exist that combine the identification and tracking of paper artifacts with the capturing of pen input and outputting information onto the paper artifacts. Our Strip'TIC application for managing paper flight strips tracks the position of paper strips using ARtags, allows users to input information by writing directly on the strips and provides feedback by projecting augmented information directly onto the paper strip. In addition, the paper strips are positioned on a back projected surface with capability for both Anoto and multi-touch input.

We next look at the types of interactions that are possible with direct input techniques such as multi-touch, pen input and tangible objects.

2.2 Interaction Techniques

Multi-touch surfaces have led to a wide range of gestures and interactions. For selection and dragging tasks on a table top surface, Forlines et al. [8] show that for a uni-manual task users performed better using indirect mouse input compared to direct touch. However, for symmetrical bi-manual tasks direct-touch was superior. In [20], gestures are described for rotating, panning and resizing documents on a multi-user interactive surface. In [28], Wu et al. present a series of gestures for multi-touch surfaces and discuss guiding principles for designing gestures. Wigdor [27] expands on this by defining a classification system for gestures based on the number of fingers involved, the number of shapes and the type of movement. Morris et al. [18] explore cooperative gestures, where the system interprets the input of multiple users as contributing to a single command.

Combining pen technology with multi-touch table top displays opens up new interactions that combine the expressiveness of touch gestures and interactions with the precision of pen technology. However research concerning interaction techniques that combine touch and pen inputs deals primarily with touch surfaces that support both touch and pen inputs as opposed to interfaces involving pen on paper input. Yee [29] investigates single touch and pen interactions, suggesting that the finger be used for panning and the pen for drawing. Brandl et al. [4] compares bimanual interactions using two pens, two handed touch and a combination of pen and touch. They find that the pen-touch combination was superior (faster and fewer errors) to either touch or pen interaction alone. They discuss general design principles for combining pen and touch input and in particular the difference between the roles of the dominant and non-dominant hand. Hinckley et al. [9] identify nine design considerations for combining pen and touch interactions and discuss how people's behaviour when interacting with physical pen and paper relates to touch+pen interactions. They propose a clear distinction between the functionalities of pen and touch: the pen writes and touch manipulates, although they encounter some situations where this rule needs to be ignored particularly when interacting with menus, users expected pen and touch to be interchangeable. For information visualization tasks, Walny et al. [26] show simi-

lar results, finding that users have clear expectations about what type of interactions should be touch or pen based.

Tangible user interfaces have many advantages over traditional graphical user interfaces [13]. They provide direct haptic feedback as the user is able to manipulate real physical objects and does not need to wait for indirect visual feedback on screen. Tangible objects provide persistence; even if computers fail they still work. They provide a seamless representation of information across physical and digital domains and they tend to encourage two-handed and multi-user interaction.

Pen and paper interfaces form a subset of tangible user interfaces. Users can manipulate paper objects by moving, rotating, stacking, folding them [11] and as well as writing and drawing on them and pointing to them with a pen [30].

Holman [11] describes basic gestures such as holding, collating, flipping, rubbing, stapling, pointing and two-handed pointing for interacting with PaperWindows which can be moved about in three dimensions. Do-Lenh et al. [7] suggest a set of finger gestures suitable for interacting with paper on a touch surface. The finger gestures may either interact with the touch surface or with the paper. As well they define gestures for transforming the paper: moving, rotating and covering.

Yeh [30] divides pen and paper interactions into two categories: drawing and commands (gestures). Within these categories, he suggests that there are four types of interactions: selection, writing, drawing, gestures.

In this paper, we extend this work by investigating how multi-touch, Anoto and augmented paper can be combined in the real world application of air traffic control at an approach control centre. We show how the combination of technologies allows us to leverage the strengths of each technology and point to new interactions that utilize multiple technologies.

The following section describes the different types of air traffic control centres followed by a description of the process of stack management, one of the tasks performed in an approach control centre.

3 Air Traffic Control Centres

Based on the portion of the flight for which they are responsible, air traffic control centres are divided into three categories: tower, approach, and en-route. The tower control centre is responsible for the plane from the moment it begins take-off until the plane is in the air and appears on the radar screen. At this point, responsibility is passed to the approach centre. The controllers at the approach centre are responsible for the aircraft until it has left the airspace around the airport. At this time, control is passed to an en-route centre. A series of en-route centres will pass control from sector to sector and possibly country to country until the aircraft approaches its destination airport. Control is then passed to the approach centre for the destination airport. The approach centre accepts responsibility for aircraft arriving from multiple directions and must schedule their arrival at the runway to ensure optimal spacing between flights. When the plane has been scheduled for landing and is on final approach to the runway, responsibility is passed to the tower control centre.

Our work focuses on the tasks performed at approach control centres. An approach centre will typically have two or three controllers depending on the volume of traffic.

- The planner controller communicates with the en-route centres and accepts responsibility for arriving aircraft. The planner is responsible for scheduling the order in which the aircraft will land.
- The radar controller communicates with the pilots and executes the plan developed by the planner controller.
- The intermediate approach controller handles the aircraft from the time they pass the final beacon until control is passed to the tower control centre. At smaller airports, this position is not used and the radar controller retains responsibility for the aircraft up until control is passed to the tower.

When the volume of air traffic is high, the planner controller, must often organize the arriving aircraft in holding stacks prior to scheduling them for landing. The process of managing these stacks is described in the following section.

3.1 Stack Management




Fig. 1. The departure of aircraft from multiple holding stacks must be coordinated to ensure optimum spacing between the aircraft.

Holding stacks are used to delay aircraft when they cannot land, typically due to congestion or weather conditions. Planes fly in horizontal loops as shown in Figure 1. A large airport typically has multiple holding stacks and the planner controller must coordinate the planes in each stack and time their departure from the stack to provide optimum spacing between the aircraft, ensuring both safety and efficiency.

Planes enter the stack at the highest level and are successively given clearance to descend to the next level. When a plane exits the stack, a known and fixed amount of time elapses before the plane lands on the runway. When two stacks are involved, these times are usually different. For example, at the Orly airport when landing east-bound on the QFU 06 runway, the time required when leaving the ODILO stack is 9 minutes and when leaving the MOLBA stack the time is 17 minutes. The air traffic controller, must mentally calculate when each plane should leave the stack, interleaving planes from each stack and maintain a 90 second gap between landings.

As shown in Figure 2, air traffic controllers organize their flight strips on the strip board to help provide a mental image of the aircraft in the holding stacks. The strips at



Fig. 2. The flight strips on the strip board are organized into three groups. The upper left and right groups contain the strips for aircraft in each of the two stacks. The centre group holds the strips for the aircraft that have left the stack and are approaching the runway beacon.

the upper left and right represent the aircraft in the two stacks and the strips in the centre represent the flights that have been cleared to exit the stack. The time of departure from the stack and the expected landing time are written on the strip.

Our augmented strip board, shown in Figure 3, maintains the same arrangement of flight strips, but provides the air traffic controller with augmented information. The flight levels within each stack are clearly shown along the left and right sides of the board and lines connect each flight strip with its level within the stack. When planes are scheduled to depart the stack, the departure time and the expected arrival time are automatically calculated and projected onto the flight strip.




Fig. 3. Augmented information clearly indicates the level of the aircraft in each stack. The times for departing the stack are automatically calculated and displayed.

We next describe our augmented flight strip board and the technical challenges of providing multi-touch capabilities.

4 Augmented Strip Board

4.1 Previous Work with Strip'TIC

Strip'TIC [12,14] was developed to bridge the gap between paper flight strips and fully electronic systems. The challenge was to provide all the efficiency and safety of an electronic system while continuing to allow air traffic controllers to interact with and write on paper flight strips. Strip'TIC was initially designed as a pen-driven interface. This allowed controllers to write on the paper strips and have the information recorded electronically. To provide feedback, we needed to track the position of the

strips. This was accomplished using AR tags on the bottom of the strips and a camera located beneath the strip board. Top projection was added to display the feedback directly on the paper strips. Bottom projection was used to display additional information on the strip board such as a virtual representation of a flight strip when the paper strip was removed from the board, and buttons for selecting different modes of operation and for inputting data. We also wanted the users to be able to interact with both the strip board and the radar screen and thus these surfaces were covered with Anoto film.

During initial user testing, it was found that air traffic controllers wanted to be able to use their fingers to interact with either the paper or virtual strips. Thus, we began investigating how to add multi-touch interactions. The next section describes some of the technical challenges in adding touch capabilities to the strip board.

4.2 Multi-Touch Technical Challenges with Strip'TIC

The physical shape of the strip board makes multi-touch interaction challenging and makes many of the standard techniques for touch interaction such as Frustrated Total Internal Reflection (FTIR), Diffuse Illumination (DI) and LED frames infeasible. Instead of a smooth surface as found on typical multitouch tables, the strip board contains a series of ridges or steps that allow the controllers to easily align the paper strips horizontally and keep them in place. To provide a light source for detecting touches, we shine a layer of infrared light above the surface of the strip board [24]. When the surface of the board is touched, this light is reflected down through the strip board and is detected by a camera beneath the board (see Figure 4). The Community Core Vision (CCV) library processes the camera image and outputs finger tracking data. To create the layer of light, a row of 48 infrared LED's are attached slightly above the strip board along the left and right sides of the board as shown in Figure 5. Because the paper strips are flat and below the level of the infrared light, they are not detected and do not interfere with touch detection. However, the surface of the strip board itself does cause multiple false touches along the edges of the ridges. Fortunately, the bands where these false touches occur are narrow and are found in areas where the user does not need to touch. Thus, in software, we are able to ignore any touches that originated within one of these bands.




Fig. 4. Light from the infrared LED's is reflected through the strip board when it is touched



Fig. 5. Infrared lights along one side of the strip board

4.3 Pen Orientation Detection

Because the surface of the strip board is covered with transparent Anoto film, users can interact interchangeably with either the Anoto pen or with touch gestures. The Anoto pen uses an infrared LED to illuminate the Anoto pattern on the strip board. Some of the light from the Anoto pen LED is transmitted through the strip board, allowing the camera used for touch detection to interpret it as a touch. Thus, we have two positional inputs from for the pen, one based on the position of the pen relative to the Anoto pattern, and the other from the touch detection. When the pen is in use, we ignore “touches” generated by the pen and use only the Anoto input. One interesting side effect of obtaining both touch and pen positions when using the Anoto pen is that it allows us to determine the orientation of the pen, and from that infer whether the pen is being held in the left or right hand of the user. As shown in Figure 6, when the pen is held in the right hand (and oriented up and to the right) the light from the pen is detected to the right of the pen tip. When the pen is held in the left hand the light from the pen is detected to the left of the pen tip.




Fig. 6. When the pen in held is the right hand the light from the pen is detected to the right of the pen tip. When the pen is held in the left hand the light is detected to the left of the pen tip.

Other research has used techniques such as hand occlusion [5] or pens with Wacom tablets [23] to determine the pen orientation. The light detection from our touch surface provides an additional technique.



Fig. 7. The strip board allows for the investigation of other input devices such as (a) a static raised area on the board for entering flight headings or (b) a device with three LED's that can be placed anywhere on the board

4.4 Tangible Input Devices

The multi-touch surface of the strip board creates the potential for other input devices. We investigated two potential devices for entering flight headings, represented as a value from 0 to 355 rounded to the nearest five degrees. The first device (Figure 7a) consisted of a raised ring of plastic mounted on the strip board. The position of the user's finger on the ring was used to calculate the heading angle. The second device (Figure 7b) consisted of three LED's forming an isosceles triangle powered by a nine-volt battery. The device could potentially be placed next to a flight strip to enter the heading for that flight.

Thus we see that in addition to providing finger detection, our touch detection system also allows for the detection of both pen orientation and of other tangible objects.

5 Interactions with Paper Flight Strips

Our observations at Toulouse Blagnac Airport approach control centre and at the air traffic control simulation centre at École Nationale de l'Aviation Civile in Toulouse, confirmed that air traffic controllers at approach centres interact with paper flight strips in a manner similar to air traffic controllers at en-route centres [16]. Air traffic controllers arrange strips using one hand or two hands, and point at strips to indicate issues with or potential conflicts between aircraft. The controllers switch rapidly and frequently between holding a pen to write and using the same hand to manipulate or point to flight strips. The pen is also used to point to locations on the strip board and on the radar screen. In the approach centre, there is a segregation of duties between the planning controller who coordinates with the en-route control centres and passes control of aircraft to and from these sectors and the radar controller who communicates with the pilots to execute the flight plan determined by the planning controller. Frequent collaboration occurs between the two controllers, with the planning controller inserting new strips on the board of the radar controller. The planning controller also points to strips to discuss issues and at times observes the strips on the radar controller's strip board.

In developing interactions with augmented flight strips, we wanted to preserve as much as possible the existing interactions with paper strips while providing opportunities for increased safety and efficiency that computer based systems have the potential to provide.

As an aircraft descends through multiple levels within the holding stack, the flight strip can become covered with circles and lines. We preserve the same interaction, allowing the controller to write on the flight strip using the Anoto pen. The system recognizes the circles and lines and highlights the last level selected using top projection. In addition, when a flight level clearance is given, a thick white line appears that connects the flight strip to the appropriate level in the stack. When the flight strip is moved on the strip board, the top projected images and the white line follow the strip.

We also created several alternate interactions for entering flight level clearance information into the system, such as buttons located along the sides of the strip board.

6 User Study

Our augmented strip board has the advantage of combining traditional paper-based and AR-based digital interaction with flight strips. We wished to address the question of whether users were capable of usefully working with all of these modalities, or whether they led to confusion.

To investigate these questions, we held design sessions with 10 air traffic controllers with 5 to 35 years experience (average 16 years) from eight different approach centres in France including Paris, Orly, Brest, Marseille and Blagnac. Seven of the participants were male and three were female; eight were right handed and two were left handed. Eight of the controllers had previously seen a demonstration of the Strip'TIC system; however, none had used the system. Nine had experience using a touch surface such as an iPad or smart phone and one participant had previously used a digital pen. Each session lasted approximately 45 minutes.

During the sessions, the system was demonstrated to the controllers who were given time to experiment with different interactions and become familiar with the system. The controller was then given a series of tasks to complete using different interaction techniques. Semi-structured interviews were used to obtain the controllers' impressions of the interaction techniques. All of the sessions were video taped. The tasks performed by the controllers were as follows:

1. Writing Commands
 - On the paper flight strip enter flight level 120
 - On the virtual flight strip enter flight level 120
2. Freehand Writing
 - On the paper flight strip indicate that there is a radio problem
 - On the virtual flight strip indicate that there is a radio problem
3. Moving Strips
 - Move a paper strip horizontally and vertically
 - Move a virtual strip horizontally and vertically using the finger

- Move a virtual strip horizontally and vertically using the pen
4. Pen vs Touch with Buttons on Left
 - Assign flight level 130 using touch only
 - Assign flight level 120 using pen and touch
 5. Pen vs Touch with Buttons on Right
 - Assign flight level 80 using touch only
 - Assign flight level 70 using pen and touch

In addition, six of the controllers completed a scenario that involved several tasks.

1. Using any technique you wish assign the following flight levels:
 - Flight: AF113LS – Level: 110
 - Flight: BMM2010 – Level: 120
 - Flight: BZ716WH – Level: 140
 - Flight: AMC466 – Level: 60
 - Flight: AF015TM – Level: 70
2. Switching Flight Levels:
 - Change AF113LS to flight level 120 and BMM2010 to flight level 110
 - Change AF015TM to flight level 60 and AMC466 to flight level 70
3. Writing: Make a note on one of the strips that the transponder has failed
4. Schedule the next three flights for leaving the stack
5. Add two more flights to the stack and assign them a suitable flight level

7 Observations

7.1 Writing Commands on Virtual or Paper Strips

All of the participants found it much easier to circle the flight level on the paper strip than on the virtual strip. With the virtual strip there was always a slight offset between where the participant saw the tip of the pen and where the projected ink marks were displayed. The offset was due to the thickness of the glass and the angle at which the participant viewed the display. The participants were consistently more accurate when entering a flight level directly on the paper.

7.2 Freehand Writing on Virtual or Paper Strips

Again, all of the participants preferred writing on the paper strip compared to the virtual strip. Although the position of the pen was not as critical, the lag in the feedback when writing on the virtual strip, made writing more difficult.

7.3 Moving Flight Strips

Moving the physical paper strips about the board was found to be the quickest way to reposition strips on the board. However eight of the ten participants, all of whom had

previous experience using a touch surface, found using touch to select and drag the virtual strips about the board nearly as simple and efficient provided that two of the strips did not become overlaid one on top of another. When this occurred, it was difficult to select the strip that they wanted to move. With the paper strips, it was easy to separate strips that were on top of each other. None of the participants felt that the pen was useful for moving the virtual strips.

7.4 Selection Tasks - Interaction Directed

We created an interaction in which the participants could assign a flight level by selecting a button near the edge of the strip board and then selecting one of the flight strips. The interaction was designed so that they could select the button and the strip in either order or simultaneously. They could also use either the pen or their finger for one or both of the touches.

In general the controllers, found that touch and pen worked well for pointing tasks such as selecting flight levels. As shown in Figure 8a, four of the ten controllers preferred using the pen, while the remaining six preferred using touch.




Fig. 8. Users indicated a slight preference for touch over pen interactions. However during the unguided scenario, the users were equally divided between using touch, using the pen and using a combination of both modalities

7.5 Selection Tasks - Task Directed

Six air traffic controllers completed the final unguided scenario in which they performed selection tasks to assign initial flight levels and to change flight levels. Of these six controllers, two used touch for all the interactions, two used the pen for all the interactions and two used a combination of pen and touch (Figure 8b).

7.6 Scheduling Leaving the Stack

All six of the controllers who completed this task moved the paper strips into the correct position to complete this task. No one attempted to remove the paper strips and move the virtual strips.

The feedback from the controllers about the stack management time calculations was uniformly very positive. All six controllers liked how simple it was to move the

flight strips and have the times calculated and displayed automatically. They felt that having the times calculated automatically would be quicker and in dense air traffic could be very helpful for the planner.

7.7 Devices for Flight Headings

The two devices for entering flight headings (Figure 7) were demonstrated to six of the air traffic controllers. They all had strong opinions about the devices, although there was no consensus. Three of the controllers liked the device with the LED's. They liked the possibility of using the device anywhere on the strip board and that it could be used in either hand. Right handed controllers did not like the idea of using their left hand to enter the headings with the ring device. One controller, Participant 10, preferred the ring that was mounted on the strip board, but did not like its current location. She felt that adding the LED device was just one too many things. It was ok to have paper and pen and touch, but she did not want another device to pick up. Two controllers did not like either option, with Participant 7 suggesting a form of keypad might be more suitable for entering the headings.

7.8 General Comments

The feedback on the system from the air traffic controllers was all very positive. They liked the visual representation of the flight levels at the sides of the strip board and thought it would be good for organizing stacks. The controllers liked the idea of retaining paper strips, especially for passing the strips to another user. They also thought that the idea of a virtual strip was useful commenting "Paper can get lost so a virtual strip underneath could help with that".

8 Discussion

Strip'TIC was initially designed as a pen driven interface. However, early user testing indicated that air traffic controllers would also like to be able to use their fingers to interact with either the paper or virtual strips. Thus, we added multi-touch capabilities. Through our user study, we attempted to assess whether this added modality was confusing rather than helpful, and to extract design guidelines for our prototype.

8.1 Combining Multiple Modalities

We used our study to address three potential concerns about combining multiple modalities:

Ease of Training: Did having multiple modes for completing a task make it difficult for the air traffic controllers to learn how to use the system?

In general, the air traffic controllers had little difficulty in learning to use the system. With less than 20 minutes of training time, they were able to use all the modali-

ties and found it straightforward to do so. Air traffic controllers are expert users of their systems, and so 20 minutes of training is a negligible cost.

Cognitive Overhead: Did having multiple modes of completing a task require the air traffic controllers to think about which mode to choose before completing the task?

Having multiple ways of completing the same task can cause users to spend more time thinking about which technique to use to complete the task rather than about the task itself. In general, we found that users quickly selected one mode for completing a task and continued to use that mode for all the tasks. We did see some examples of indecision, for example, when first beginning the task involving assigning flight levels to five flight strips, participant 9 initially reached for the pen. However, he did not pick it up and then used touch interactions to enter all the flight levels.

Confusion: Were the air traffic controllers confused by the limitations of some of the modalities?

Confusion may arise when an interaction does not work in the same manner on all system features. The current implementation of Strip'TIC has several potential issues of this type:

- The pressure of writing with a pen on a paper strip can cause the strip to move, requiring the user to hold the strip in place with the non-dominant hand when writing. However the virtual strip does not move when writing in the centre area of the strip. In contrast, placing a hand on a paper strip will prevent it from moving while writing, but placing a hand on the virtual strip will make the strip move, which is not expected.
- Users can touch a virtual strip to select it; however, because the camera used for touch detection is located under the strip board, the camera could not “see” touches on the paper strips. Thus it was not possible for users to select a flight strip by touching the paper strip.
- Depending on the location of the buttons for selecting flight levels, there was a difference in the ease of using the pen for assigning flight levels. For a right-handed user, it was easy to select a flight strip using the pen and then use their left hand to select a flight level to the left of the strip. However, if a person used the same approach for selecting a flight level to the right of the strip, he would become cross-handed and the interaction would be exceedingly awkward.

We saw instances of confusion when the participants first used Strip'TIC. For example, participant 10 tried placing her non-dominant hand on the virtual strip to hold it in place when writing on it and participant 1 became cross-handed when using a combination of pen and touch to enter flight levels on the right hand side of the strip board. However, these users quickly adapted and did not display any signs of confusion upon subsequent interactions. We found that all six of the air traffic controllers who completed the final unguided scenario adopted essentially the same workflow and used all of the modalities: pen for writing on strips, touch for selecting strips and flight levels and paper for arranging the flight strips on the strips board. The availabil-

ity of different modalities did not overburden the users and they did not find it difficult to determine the appropriate modality of each interaction.

8.2 Technical Considerations

The Strip'TIC system addresses a range of complex challenges: combining paper flight strips, digital pen and finger tracking. Our implementation of touch interaction is simple, flexible and allows finger detection, pen orientation and tangible object detection.

Our user study highlighted some of the technical limitations of the system primarily related to the use of Anoto pen technology as well as areas for future implementations and investigations. These are discussed in the following sections.

Anoto Pen: Our results show that Anoto pens, while suitable for written input on paper surfaces, should not be used for written input on digital surfaces. There are four technical issues that need to be resolved:

- *Accuracy:* Due to a combination of parallax and alignment of the bottom projected image with the Anoto film, there was a slight difference between where the participant saw the tip of the pen and where the tip was detected on the virtual flight strip that made writing on the virtual strip difficult.
- *Feedback time:* When writing on the paper strips, the controllers received instant feedback seeing the physical ink on the paper. However, on the virtual strip, the delay between moving the pen and seeing the virtual ink trail made writing difficult. Surprisingly, the air traffic controllers found the lag more problematic for freehand writing than for circling flight levels.
- *Pen Angle:* The angle of the pen was also an issue when writing on the virtual strips. On paper, the pen functioned well as long as it was held at 35 degrees or more above horizontal. However on the strip board, the pen needed to be held at an angle greater than 45 degrees. Half of the participants initially held the pen at an angle between 35 and 45 degrees and had to adjust their grip of the pen in order to write on the strip board.
- *Feel:* The contact between the pen and glass was not as pleasing as between the pen and paper.

The Anoto pens were also not suitable for dragging virtual strips on the strip board. The issues here were due to the design of the strip board and of the virtual strip. First, because the pen could be used for writing on the strip as well as for moving the strip, we designated only the 0.5 cm coloured border at the top of the strip as the region that needed to be selected in order to move the strip. The remainder of the strip was reserved for writing on the strip. The users found that this border at the top of the strip was too narrow to select accurately and they often began writing instead of moving. Second, due to the stepped surface of the strip board, the pen often lost contact with the board when trying to move the strips vertically, requiring the users to try repeatedly before the strip would move up or down to the next row.

The Anoto pens did work well for selection tasks on both the paper flight strips and on the strip board.

Touch Detection: We found that multi-touch is a natural feature for moving paper and virtual strips and that users valued having this additional modality of interaction. However, our current implementation is based on a camera located underneath the strip board and cannot detect when a user touches a paper strip. Thus it is not possible for paper flight strips to be selected by pointing at or touching the strip. We plan to investigate whether other technology such as Microsoft’s Kinect or a LED light frame might enable this interaction.

Our current touch detection system, does allow for the detection of both pen orientation and tangible input devices. To date, we have not fully explored the range of possibilities afforded by these inputs. Future work will investigate other possible input devices and also how the system can adapt to different users, based on the pen orientation.

9 Conclusion

In this paper, we have described our design evolution with the Strip’TIC project. Our technical contribution addresses the complex challenges of combining multi-touch, digital pens and tangible objects. However, it was not obvious whether combining these technologies necessarily leads to a better user experience, and whether more options and modalities would be helpful or merely lead to confusion.

Through our user study we found that users adapted quickly to the interface. We found that all the controllers who participated in an unguided interaction scenario adopted essentially the same workflow and used all of the modalities: pen for writing on strips, touch for selecting strips and flight levels and paper for arranging the flight strips on the strips board. The availability of different modalities did not overburden the users and they did not find it difficult to determine the appropriate modality to use for each interaction.

References

1. Aitenbichler, E., and Schnelle-Walka, D. An extensible architecture for multitouch & pen interactive tables. Proceedings of EICS (2010).
2. Anoto, digital pen technology. <http://www.anoto.com>.
3. ARToolkit. <http://www.hitl.washington.edu/artoolkit>.
4. Brandl, P., Forlines, C., Wigdor, D., Haller, M., and Shen, C. Combining and measuring the benefits of bimanual pen and direct-touch interaction on horizontal interfaces. In Proc. AVI, ACM (2008), 154–161.
5. Brandl, P., Leitner, J., Seifried, T., Haller, M., Doray, B., and To, P. Occlusion aware menu design for digital tabletops. In Proc. CHI, ACM (2009), 3223–3228.
6. Cohen, P., and McGee, D. Tangible multimodal interfaces for safety-critical applications. Communications of the ACM 47, 1 (2004), 41–46.
7. Do-Lenh, S., Kaplan, F., Sharma, A., and Dillenbourg, P. Multi-finger interactions with papers on augmented tabletops. In Proc. TEI, ACM (2009), 267–274.

8. Forlines, C., Wigdor, D., Shen, C., and Balakrishnan, R. Direct-touch vs. mouse input for tabletop displays. In Proc. CHI, ACM (2007), 647–656.
9. Hinckley, K., Yatani, K., Pahud, M., Coddington, N., Rodenhouse, J., Wilson, A., Benko, H., and Buxton, B. Pen+ touch= new tools. In Proc. UIST, ACM (2010), 27–36.
10. Hofer, R., and Kunz, A. Digisketch: taming anoto technology on lcds. In Proc. 2nd ACM SIGCHI symposium Engineering interactive computing systems, ACM (2010), 103–108.
11. Holman, D., Vertegaal, R., Altosaar, M., Troje, N., and Johns, D. Paper windows: interaction techniques for digital paper. In Proc. SIGCHI, ACM (2005), 591–599.
12. Hurter, C., Lesbordes, R., Letondal, C., Vinot, J., and Conversy, S. Strip’TIC: exploring augmented paper strips for air traffic controllers. In Proc. AVI, ACM (2012), 225–232.
13. Ishii, H. Tangible bits: beyond pixels. In Proc. TEL, ACM (2008), xv–xxv.
14. Letondal, C., Hurter, C., Vinot, J., Lesbordes, R., and Conversy, S. Strip’TIC: Designing a paper-based tangible interactive space for air traffic controllers. In Proc. CHI, ACM (2013).
15. Liwicki, M., Rostanin, O., El-Neklawy, S., and Dengel, A. Touch & write-a multi-touch table with pen-input. 9th Int. Workshop on Document Analysis Systems (2010), 479–484.
16. MacKay, W. Is paper safer? The role of paper flight strips in air traffic control. ACM Transactions on Computer-Human Interaction (TOCHI) 6, 4 (1999), 311–340.
17. Mackay, W., Pothier, G., Letondal, C., Bøegh, K., and Sørensen, H. The missing link: augmenting biology laboratory notebooks. In Proc. UIST, ACM (2002), 41–50.
18. Morris, M., Huang, A., Paepcke, A., and Winograd, T. Cooperative gestures: multi-user gestural interactions for co-located groupware. In Proc. CHI, ACM (2006), 1201–1210.
19. Paelke, V., and Sester, M. Augmented paper maps: Exploring the design space of a mixed reality system. *Journal of Photogrammetry and Remote Sensing* 65, 3 (2010), 256–265.
20. Ringel, M., Ryall, K., Shen, C., Forlines, C., and Vernier, F. Release, relocate, reorient, resize: fluid techniques for document sharing on multi-user interactive tables. In Proc. CHI, ACM (2004), 1441–1444.
21. Song, H., Guimbretiere, F., Grossman, T., and Fitzmaurice, G. Mouselight: bi-manual interactions on digital paper using a pen and a spatially-aware mobile projector. In Proc. CHI, ACM (2010), 2451–2460.
22. Steimle, J. Survey of pen-and-paper computing. In *Pen-and-Paper User Interfaces, Human Computer Interaction Series*. Springer Berlin Heidelberg, 2012, 19–65.
23. Taub, D. The BoPen: a tangible pointer tracked in six degrees of freedom. PhD thesis, Massachusetts Institute of Technology, 2009.
24. Teiche, A., Rai, A., Yanc, C., Moore, C., Solms, D., Cetin, G., Riggio, J., Ramseyer, N., Dintino, P., and Muller, L. Multi-touch technologies, 2009. NUI Group.
25. Tsandilas, T. Interpreting strokes on paper with a mobile assistant. In Proc. UIST (2012).
26. Walny, J., Lee, B., Johns, P., Riche, N., and Carpendale, S. Understanding pen and touch interaction for data exploration on interactive whiteboards. *IEEE Transactions on Visualization and Computer Graphics* (2012).
27. Wigdor, D. Architecting next-generation user interfaces. In Proc. AVI, ACM (2010), 16–22.
28. Wu, M., Shen, C., Ryall, K., Forlines, C., and Balakrishnan, R. Gesture registration, relaxation, and reuse for multi-point direct-touch surfaces. In *TableTop*, IEEE (2006), 183–190.
29. Yee, K. Two-handed interaction on a tablet display. In Proc. CHI, ACM (2004), 1493–1496.
30. Yeh, R., Paepcke, A., and Klemmer, S. Iterative design and evaluation of an event architecture for pen-and-paper interfaces. In Proc. UIST, ACM (2008), 111–120.