Combining Point Force Haptic and Pneumatic Tactile Displays

Yeongmi Kim*
HMCI Lab,
Gwangju Institute of
Science and Technology.
www.dyconlab.gist.ac.kr

Ian Oakley†
Electronics and
Telecommunications
Research Institute (ETRI).
www.etri.re.kr

Jeha Ryu‡
HMCI Lab,
Gwangju Institute of
Science and Technology.
www.dyconlab.gist.ac.kr

ABSTRACT

The ability to combine tactile, on the skin, sensory cues with the kinesthetic ones that widely available force feedback devices are able to produce is a desirable one, opening the door to the production of more realistic, compelling virtual environments. Pneumatic air jet displays can be easily mounted on existing force feedback devices and we believe have the potential to provide useful complimentary tactile information. However, there is little concrete psychophysical data relating to pneumatic displays, a fact which hinders their adoption. This paper addresses this challenge, and presents brief psychophysical studies examining localization rate, the two point threshold, stimulus intensity and the temporal threshold of cues produced by pneumatic air jets. Using insights gained from these studies, we also present a more concrete application focused investigation where we evaluate the effect of combining tactile and force feedback cues in a simple object manipulation task in a virtual environment. We show that task completion times are significantly improved with the addition of tactile information, validating our ideas and suggesting this topic warrants further attention.

Keywords: pneumatic tactile display, psychophysics.

1 INTRODUCTION

Widely available point force haptic devices (such as the PHANToM [8]) provide high quality kinesthetic cues and are capable of displaying a range of haptic, or touch information about virtual or simulated environments. Through motors, or other mechanical elements attached to carefully designed exoskeletal structures, they are capable of rendering the shape, compliance, friction and crudely, the fine grain details and surface texture of computer generated objects. However, a generally accepted weakness of these devices is that they are incapable of presenting information directly to the surface of the skin. In order to operate a point force device a user must hold, in one way or another, a tool like end effector, and experience the forces generated through this mediating implement. Typical end-effectors take the form of pen-like tools [e.g. 15], graspable spheres [e.g. 9] or thimbles into which one's finger is placed in order to feel the forces [e.g. 8]. The physical aspects of this grasping behavior determine the tactile, on-the-skin, experience of the displayed haptic cues.

While this kind of tool-based interaction is sufficient for many situations, and indeed in some cases (such as laparoscopic surgical simulations [18]) it is an entirely desirable, it can be a problem in others. One significant example of this is where it is the intention to present a haptic experience as if the user was unencumbered by equipment; to create the illusion that they are exploring the virtual world with their bare fingers and hands, and not through a mediating tool. In this interaction style, essentially enabled by using the thimble-like end effectors mentioned above, it is desirable that users should experience not only the net reaction forces derived from touching objects in the virtual world, but also the deformations to, and changes to the pressure exerted on, the skin of their finger. The fact that this latter form of sensory cue is not conveyed with current technologies reduces the realism of the experiences they are capable of imparting, and may reduce the effectiveness with which users are able to interact with virtual worlds. Certainly, studies examining the effect of gloves (which impair cutaneous perception) on the completion of physical tasks have found them to lead to substantial reductions in performance [4].

A number of previous authors have highlighted this issue in the VR domain, and attempts to address it have focused, perhaps unsurprisingly, on creating small fingertip displays that can be used in conjunction with force feedback devices. For example, Wagner *et al.* [17] integrated a tactile pin array into their WAM (Whole-Arm Manipulator) force feedback device. They describe a study using this system that investigates a user's ability to discriminate the compliance of virtual objects and found that the inclusion of tactile information led to increased levels of sensitivity.

Debus *et al.* [5] describe a similar investigation of a combined force feedback and tactile device. They embedded a vibrating element within the handle of their force feedback hardware, essentially augmenting it with the ability to produce rudimentary tactile cues. They then investigated user performance in a simple teleoperation task when information in different sensory modalities was available to participants. They conclude that optimal performance was attained when users experienced a combination of visual, force feedback and tactile cues.

In particular, one technology we believe deserves closer attention in this research domain is that of pneumatic air-jet displays. These are devices which produce tactile sensations by blowing jets of air on a user's skin. In this our motivations largely mirror those of Amemiya and Tanaka [1] in their description of a pneumatic display designed to be mounted on the base of the finger pad. Perhaps the most significant reason why air-jet displays are well suited to integration with a force-feedback device is because they do not require a mechanical assembly at the point of contact with the user's skin. Flexible, lightweight tubing can connect the point at which the air is impelled with force with the stimulus display locus under the user's fingertip. Furthermore, with a carefully designed system it is possible that this required tubing could be conveniently routed along the existing armature of the haptic device to allow the construction of a small, lightweight, dense fingertip display. However, this mechanical simplicity has a number of attendant disadvantages. One of these is that the length (and diameter) of the tubing used clearly effects the time between the initial display of a stimulus and it reaching the user. In some situations this kind of latency may be simply unacceptable.

Furthermore, skin stimulation through air-jets has not been studied as extensively as, for example, pressure or vibration. So while we can state with confidence the localization rates and

*email: kym@gist.ac.kr
†email: ian@whereveriam.org
‡email: ryu@gist.ac.kr

two-point limits for pressure stimuli delivered by pins pressed against the skin, it remains somewhat unclear what psychophysical parameters apply to the perception of streams of air delivered to the skin. While skin-deformation definitely occurs, it is unlikely to do so in so explicable a manner as with the application of a directly quantifiable physical stimulus (in the form of a pin pushed upwards with a certain force, or to a certain position). If the skin is placed tightly against the surface of the pneumatic display, then it is likely that a corresponding feeling of pressure will result. However, if the skin is not pressed tightly enough, some air will escape, altering the sensation, and if the skin is slightly raised then the sensation will likely differ considerably.

These concerns are echoed by Amemiya and Tanaka [1], who briefly describe a complex study that attempts to determine the two point threshold for their pneumatic device and involves the variables of body site (thumb or forefinger), delivered air-pressure, air-jet nozzle diameter, and inter-jet spacing. Unfortunately, due to its complexity, it is hard to draw concrete conclusions from this study, beyond the simple observation that each variable appears to be capable of exerting a substantial effect on perceptual performance. Reflecting the effective absence of literature and the potential we see for such displays, this paper takes up this research challenge. We describe the design of a multi-element pneumatic display, and a battery of psychophysical tests intended to serve as a practical guide for stimulus design for fingertip air-jet systems. As with Amemiya and Tanaka [1], the focus of our interest in these studies is not on the dynamics of individual jet performance, but instead on the perceptual practicalities of designing an array consisting of multiple air jets; the basic questions of how such an array should be laid out on the skin to ensure that it appropriately matches human perceptual abilities. We conclude with an application focused study integrating force-feedback and pneumatic stimuli, and speculate about the future of this work.

2 PNEUMATIC DISPLAY HARDWARE

2.1 Design of Pneumatic Display

For the purposes of these studies, we constructed a simple prototype pneumatic array consisting of 25 individual air jets arranged in a 5 by 5 grid. The array was a resin based pad designed to be mounted on the base of the index finger and held in place by an elastic sheath attached to its sides and tip and running over top of the finger. It was designed to snugly fit around the entire finger and this served to mechanically limit the variability of possible placements of a user's fingers against the array surface. It featured air nozzles with an external diameter of 2.4mm (and an internal diameter of 1.5mm). In each row of the array (running across the short axis of the finger pad) the air jets were mounted directly against one another, yielding an inter-stimulus distance of 2.4mm. However, small structural supports with a width of 0.8mm were placed between the columns of the array (running down the long axis of the finger pad) resulting in a distance of 3.2mm between pairs of adjacent jets. The overall array size was 12mm by 15.2mm. Figure 1 illustrates our device.

2.2 System Architecture

Our system possessed a relatively simple configuration. A PC used an RS232 serial connection to communicate with a dedicated Mexx ATMega 128 microprocessor which controlled the state of the 25 valves which formed the array. We used a similar communications protocol to that adopted by other authors investigating tactile arrays [e.g. 3, 10] and, in line with their systems, chose to update the array state relatively rapidly. We selected a rate of 500 times a second as it enables a reasonable level of temporal precision while remaining comfortably within the bandwidth of the RS232

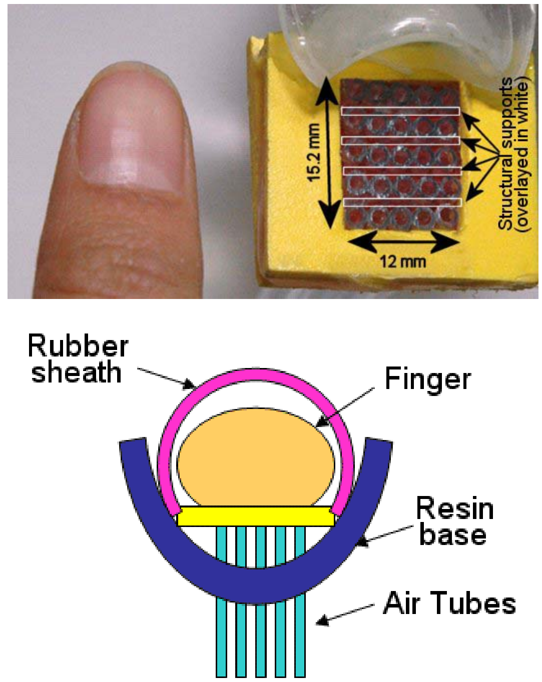

Figure 1. Photograph and schematic of initial prototype of pneumatic array.

communications link. Please refer to Oakley *et al.* [10] for more information on the design and implementation of this communication protocol. In all the studies described here we used an industrial air accumulator and regulator which ensured that the air pressure presented to users was a constant 1.034 bar. The valves used to control the flow of air were Yonwoo Pneumatic YSV10s. These are binary devices (either on or off) with a switching latency of approximately 20 ms.

3 PSYCHOPHYSICAL EXPERIMENTAL OVERVIEW

In order to gain an understanding of the kinds of tactile cue that can be effectively displayed on our pneumatic hardware platform, we engaged in a battery of brief tests. Although conducted with the rigor of psychophysics in mind, the overall goal of these studies was more focused on the rapid production of reasonable approximations than on the adoption of impeccable procedures. Consequently, a number of experimental liberties were taken, as reported in the experimental procedures described below. We are aware of the limitations such an approach conveys to the validity of our data, and believe, in the absence of a suitable substitute, that it remains a useful and valuable contribution. We would liken the approach we adopt to that taken by Tan *et al.* [12] in their informative description of the role human factors play in the design of force-reflecting interfaces.

3.1 Experiments and Participants

Two groups of subjects were used in these studies, subsequently termed groups A and B. We chose to use two groups for purely logistical reasons – it lent us the flexibility (with respect to timetabling and participant fatigue) to conduct a range of studies rapidly. Both participant groups featured eight members, four men and four women. The average age of the participants in each group was 22 and 29 respectively. Group A completed four brief experiments, two investigating localization performance and two examining the two-point threshold. Group B completed a study looking at the perception of stimulus intensity and one investigating temporal resolution.

3.2 Localization studies

The two localization studies shared a similar interface and procedure, and could be differentiated solely by the distance between the stimuli sites used in each. The intention was to contrast user performance with these different levels of inter-stimulus spacing. Each study included the presentation of 9 stimuli in the form of individual air jets and arranged in a square pattern. In the first study, these stimuli were formed by the 9 jets in the central 3 by 3 square of the 5 by 5 array (subsequently termed the dense study). In the second study, 8 of the stimuli fell on the rim of the array (the four corners and the four mid-points between them) and the 9th remained in the center (subsequently termed the spread study). This is illustrated in Figure 2. The horizontal and vertical distances between stimuli centers in the dense study were 2.4mm and 3.2mm respectively. In the spread study these values were 4.8mm and 6.4mm. In each study, each stimulus was presented a total of 20 times (leading to a total of 180 trials) and in a random order.

Each trial commenced with a screen instructing participants to press a key to begin. Upon completion of this action there was a 1 second pause, followed by 500 ms of stimulus presentation. Participants then had to press a key on the numeric keypad (the square arrangement typically situated on the right of a keyboard) to indicate the location of the displayed stimuli. The bottom left of this keyboard (the key marked with the number 1) corresponded to a stimuli on the bottom left of the fingertip array. Similarly, the number 9, at the top right of the keypad, indicated a stimuli at the top right of the fingertip. The graphical interface to the study matched the spatial layout of the numeric keypad (and featured appropriate numbering) and after each trial graphical highlighting took place to indicate both the user's response and the correct answer. After this stage, a new trial began. Prior to each experimental session participants completed a practice session which was identical in structure, but consisted of only half the number of trials.

3.2.1 Localization Results

The localization rates for each air-jet in each study are presented in Figure 3. ANOVAs revealed significant differences in these data in the case of the spread study ($F(8, 7) = 2.247$, $p<0.05$), but not in the case of the dense study ($F(8, 7) = 2.01$, $p=0.06$). Post-hoc t-tests showed the only significant difference in the spread study was between the air-jets one and nine. A t-test comparing localization performance between the two studies showed a significant difference ($p<0.001$).

Generally speaking, these results indicate that perception of air-jet stimuli from our array is relatively homogenous across the finger pad. This serves more to confirm the usefulness of our simple hardware design than offer new psychophysical insight. One caveat is the significant drop in performance observed in one of the air-jets positioned on the extremity of the array in the spread study. A likely interpretation of this result is that this air-jet was located sufficiently far under the curved edge of the finger that the gap between its outlet and the skin of the finger became large enough to impair perception in some participants. Consequently, we suggest that to ensure consistent perception it may be advisable to use an array with smaller overall dimensions than that employed in our spread study. However, this is likely to be a tradeoff, as the jets in such an array would have to spaced more closely together, potentially reducing user performance (as seen from the results achieved in our dense study). An alternative solution might be to design a more sophisticated array that is curved to better fit the contours of the finger.

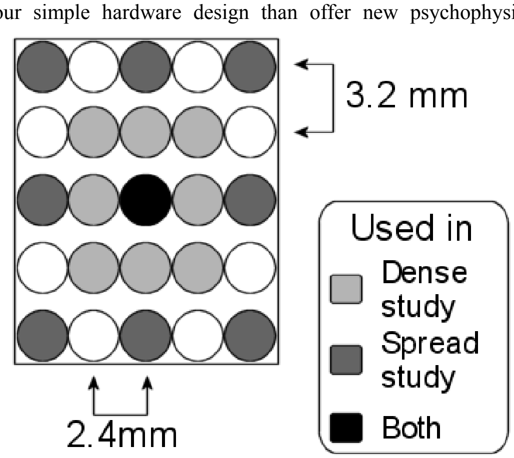

Figure 2. Diagram of pneumatic array showing stimuli used in localisation studies.

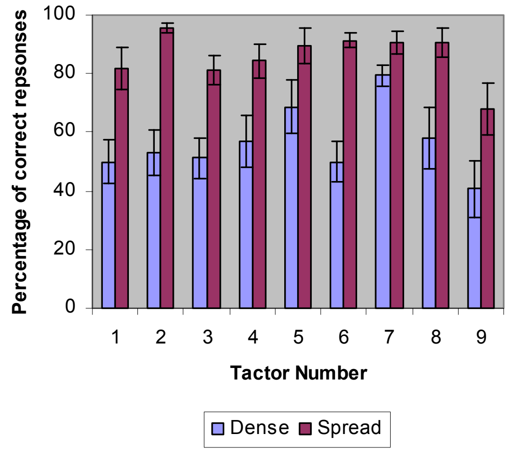

Figure 3. Data from localisation studies.

Relating these results to prior finger pad localization studies, typically conducted with pressure stimuli generated by pressing pins against the skin and in which localization rates of as low as 0.15mm have been reported [7], we can conclude that air-jet cues are not perceived with the same high levels of accuracy. In our dense study, featuring stimulus sites several mm apart, the error rate hovered around the 50 percent mark, suggesting users experienced considerable difficulty with the task. One possible reason for this is that the feel of the air-jet stimuli may be dependant on the position of the user's finger pad on the array. Small movements laterally or, perhaps more significantly, vertically away from device may cause considerable variations to the cues, and are challenging to measure or control for. An alternative explanation lies with the size of the pneumatic cues. The inner diameter of the tubing we used was 1.5mm, considerably larger than the point of a pin, and yielding effective inter-stimulus spacing (measured from the extremities of the air nozzles) of as little as 0.9mm. It is also possible that upon exiting the tubes, the jets of air immediately began to spread out causing a still larger stimulus footprint on the skin. These issues remain unresolved at this time and warrant further investigation.

3.3 Two-point studies

Two studies were conducted to gauge the two point threshold for air-jet stimuli. They were intended to complement one another, with the second examining smaller scale inter-site distances than the first. Both experiments shared a similar task and procedure. In each, participants tapped a key to begin a trial, and after a 1 second pause were presented with a stimuli which they then had to judge as either consisting of one or two separate jets of air. Feedback was given regarding their response but not as to its correctness. Both studies were preceded by practice sessions with half the duration of the experimental sessions.

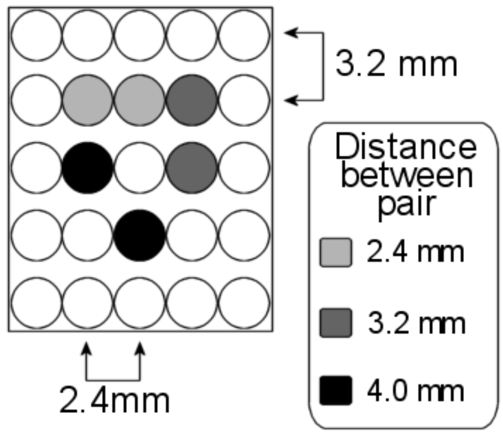

Figure 4. Diagram of pneumatic array showing sample set of stimuli used in second two point threshold study.

The first study consisted of 5 different stimuli: a single jet, and four pairs of two jets. All the stimuli were situated on the centre column of the array (positioned along the centre of the long axis of the finger). The stimuli composed of pairs of jets featured either adjacent jets, or those separated by 1, 2 or 3 spaces. This corresponds to distances (as measured from the centers of the jets) of 3.2, 6.4, 9.6 and 12.8 mm. Each of the 5 stimuli was presented 20 times (leading to a total of 100 trials), and the position of each in the central column was randomized within the physical limits of the design (for instance, there are 5 possible display sites for the stimuli consisting of a single jet, but only 1 for the stimuli consisting of 2 jets separated by 3 spaces).

The second two-point study examined a set of somewhat closer points, and took advantage of the physical constraints of our array. It consisted of four stimuli, one of which was generated by a single air jet, the remaining three generated by a pair of jets. The pairs of stimuli were all adjacent, but differed in the directionality of this adjacency. As our array is not uniformly spaced along its axes (as illustrated in Figure 1, and due to the presence of structural supports between its rows, but not its columns), it features different inter-jet spacing between horizontally, vertically and diagonally adjacent jets and this fact was used to generate stimuli pairs with centers which were 2.4, 3.2 and 4.0 mm apart. This is illustrated in Figure 4. To control for possible response biases in this study, each of the stimuli pairs was presented 30 times, while the single jet stimulus was presented a total of 90 times. This equalized the number of times participants were exposed to individual and pair stimuli, and led to a total of 180 trials. All stimuli were presented on the central 3 by 3 portion of the array, and randomized for the physical limits of this configuration. This led to 9 possible locations on which to display the single jet, 6 for each of the horizontal or vertical pairs, and 8 for the diagonal pair.

3.3.1 Two-point Results

Figures 5 and 6 illustrate the accuracy with which participants correctly judged the stimulus in each trial as being composed of one or two points. As we were attempting to determine the two point threshold we discarded the data related to stimuli generated by a single point from our formal analyses. From the remaining data, ANOVA's showed significant effects (at $F(3, 7) = 128.751$, $p<0.001$ for the first, larger scale study, and $F(2, 7) = 11.052$, $p<0.001$ for the second). Post-hoc t-tests (with Bonferroni confidence interval adjustments) on the data from the first study revealed that when two points were adjacent performance was considerably lower than with any of the other configurations, where accuracy rates approached or exceeded 90 percent (all at $p<0.01$). Similar t-tests on the data from the second study confirmed significant differences between the most distantly separated pair of cues and both of the other two pairs (at $p<0.05$).

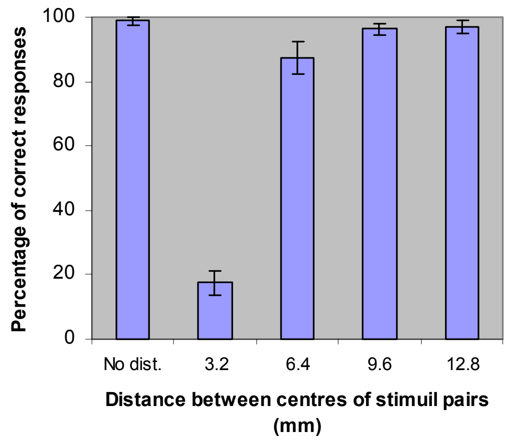

Figure 5. Data from first two-point threshold study.

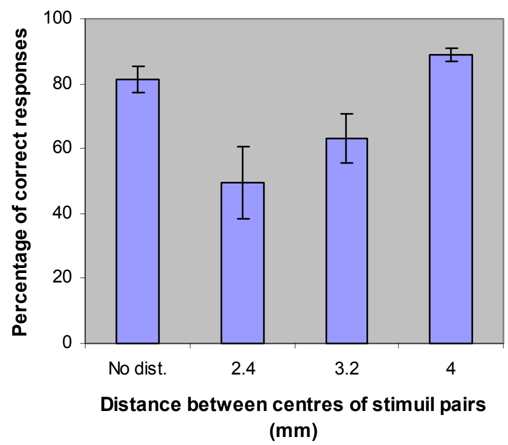

Figure 6. Data from second two point threshold study.

It is worth briefly discussing the differences in accuracy rates between the two studies. In the first study, subjects achieved greater levels of accuracy when presented with a single stimulus and much worse when presented with two reasonably proximal (3.2 mm apart) stimuli than they did in the second study (where they performed well below chance). The reason for this disparity is probably due to a response bias in the design of the first study: as subjects were presented with two response options to each trial, they most likely attempted to choose each one approximately 50% of the time. Needless to say, as the presence of a single cue occurred in only 20% of the trials, the data is skewed to reflect the expectations of the participants. Half the trials in the second study consisted of a single stimulus, avoiding this bias, and therefore providing a more accurate measure of the two-point threshold for pneumatic cues generated by air-jets.

The two-point threshold for pressure stimuli generated by pin-pricks to the skin has been reported to be as low as 1mm [6]. The data from our second study indicates that when the pneumatic cues are positioned 2.4mm apart users achieve an accuracy rating of approximately 50 percent, effectively equivalent to chance. Although somewhat better with a 3.2mm cue separation, participants were only able to reliably determine the presence of two separate cues when they featured a 4mm gap between them. This result is in line with

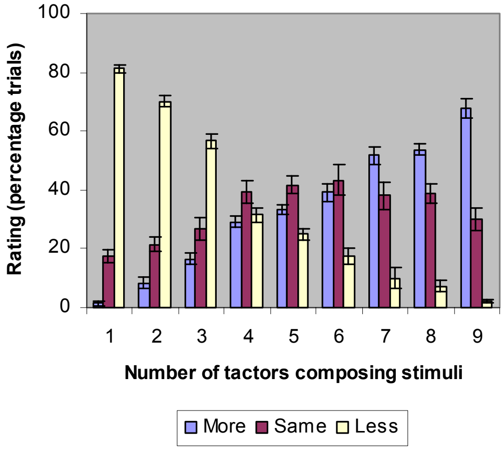

Figure 7. Ratings data from intensity study.

Number of air jets in other stimuli \ Number of air jets in one stimuli	2	3	4	5	6	7	8	9
1	0.08	0.01	0.01	0.01	0.01	0.01	0.01	0.01
2		0.46	0.01	0.01	0.01	0.01	0.01	0.01
3			0.01	0.01	0.01	0.01	0.01	0.01
4				1.0	0.64	0.01	0.01	0.01
5					1.0	0.01	0.01	0.01
6						0.07	0.14	0.01
7							1.0	0.12
8								0.20

Table 1. P vals from t-tests in intensity study comparing how often each stimuli is rated as more intense than each other (greyed cells indicate significant results).

those of our localization studies, supporting that data and further suggesting that the sensory thresholds for pneumatic cues are considerably greater than those for the more commonly studied pressure cues.

3.4 Stimulus Intensity

Unlike many forms of tactile display, one current technological limitation of most air-jet displays is that they are unable to present stimuli of different intensities. This is largely due to the fact that most valves (which are used to control the airflow), and especially those that are small, reasonably priced and quiet, are binary. They are either off or on and consequently so are the stimuli that they can create. However, delivering cues that have the capability to grow or shrink in magnitude remains an attractive goal. One way this can be achieved with a display composed of binary elements is by changing the size of the skin area to which a stimulus is applied by changing the number of array elements used to display it. At threshold levels, a relationship between the size of skin contactor and the perceived magnitude of a displayed vibrotactile cue has been established by Verrillo [16], more recently investigated by Oakley *et al.* [10] and similar effects are well documented in other sensory modalities, such as vision [13].

Reflecting this work, we conducted a study intending to gauge whether participants can effectively discriminate between stimuli composed of different numbers of simultaneously active air jets. In fact, given our previous localization and two-point threshold results, it is reasonable to suggest that by activating groups of adjacent air-jets, we can in fact stimulate different sized areas of the skin. Based on this perceptual assumption we conducted a study in which participants experienced a pair of stimuli, each created by activating between 1 and 9 air jets, and then had to judge which of these two was of greater magnitude, or if they were the identical. The air-jets used in this study were drawn from the central 3 by 3 portion of the array and we twice compared each of the 9 possible magnitudes against all this entire set. This led to a total of 162 trials (9 magnitudes times 9 magnitudes times 2 presentations). Each stimulus was always produced by a single arrangement of one or more adjacent tactors (unlike in the studies on two-point localization which explicitly varied which jets stimuli were displayed on).

In this study, each trial consisted of a subject depressing a key to begin, a 1000 ms pause, followed by a 500 ms stimulus presentation, another 1000 ms pause and a final 500 ms stimulus presentation. Participants were required to respond by pressing keys corresponding to whether they thought the first stimulus was greater than the second, that they were the same, or that the second was greater than the first. An on-screen interface reinforced these instructions, and highlighted user responses, but not their correctness. Immediately prior to the experimental session, participants completed an 81 trial practice session.

3.4.1 Stimulus Intensity results

The data recorded in this study consisted of pair wise judgments comparing the perceived intensity of each individual stimulus against the full set of stimuli. The simplest way to analyze data in this form is to simply tally the number of times a given stimulus was rated as being greater, the same or less than another. As perceived intensity increases or decreases, these three counts should similarly increase or decrease. These data are shown in Figure 7. We confined our analyses to the number of times each stimulus was rated as greater than another. An ANOVA indicated the variations in this statistic attained significance ($F(8, 7) = 129.65$, $p<0.001$), while post-hoc t-tests incorporating Bonferroni confidence interval adjustments revealed a large number of significant differences, summarized in Table 1. These results strongly suggest that by varying the size of a pneumatically created tactile stimulus, we can effectively vary its perceived magnitude. This finding is made especially valuable in light of the fact that the individual jets within our array are only capable of producing a stimulus of a single magnitude. By demonstrating the viability of an alternative mechanism to vary the intensity of a cue, we open the door to the production of a whole new range of dynamically changing stimuli.

3.5 Temporal Resolution

An important aspect to consider in the design of any display is the temporal resolution of the relevant human perceptual system. Such knowledge informs us about the quality of the stimuli we can effectively display; we need only be capable of displaying stimuli at a speed which they can be easily perceived and discriminated. With regards to the human tactile system, temporal resolution is a complex issue. Due to the presence of different types of mechanoreceptor in the skin (each of which responds best to different types and frequencies of stimuli, up to peak sensitivities of 250 Hz for vibration) it is hard to use the existing literature to determine what the appropriate temporal resolution will be for pneumatic cues. This issue is further complicated by the fact that the kind of pneumatic display we are considering in this paper exhibits a rather complex, and not fully classified, temporal behavior. Not only do the valves that control the airflow have a relatively substantial latency (in the order of 20 ms), but there is also the

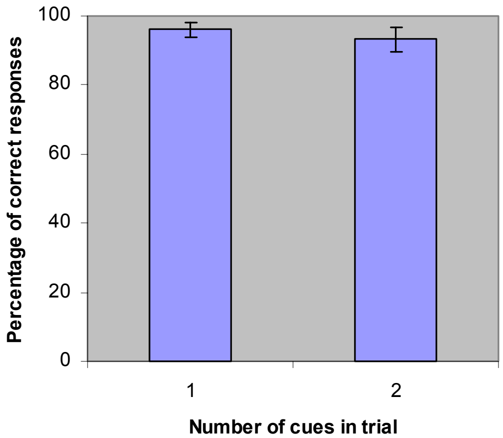

Figure 8. Ratings data from temporal study.

issue of the physical delay between the activation of a valve and the corresponding arrival of a stimulus on the skin (due to the simple fact that the air is routed down lengths of tubing).

To investigate how these factors impact on the kinds of cues we are capable of creating, we designed a study examining the perception of pairs of temporally proximate stimuli. The goal of this investigation was to (within the limits of our display hardware) determine the perceptual threshold relating to the temporal spacing between pneumatic cues. Each trial in the study involved the display of either a single cue or a pair of consecutive cues, and participants were required to press a key to make a judgment as to which of these eventualities they felt had taken place. They received graphical highlighting about their choice but not about its accuracy. The total duration of each trial was kept at a constant 100 ms. 50 trials were administered in total. Half were composed of a single stimulus, the other half by 2 stimuli separated by 2ms, the minimum duration that our software is capable of presenting and considerably under the mechanical latency of the valves used in our hardware. All stimuli were displayed on the central jet of the array. Prior to the experimental session, participants completed a short practice session with a similar structure but composed of only 10 trials.

3.5.1 Temporal Resolution Results

The results from this study are shown in Figure 8. The accuracy rate approaches one hundred percent for stimuli in both conditions, and a t-test indicated there was no significant difference between them ($p = 0.43$). From this result we can conclude that the temporal abilities of the human perceptual system considerably outperforms that of our display hardware, and that we can not expect to be able to effectively present high frequency information.

3.6 General Discussion

The motivation underlying these psychophysical tests was to attempt to determine the perceptual qualities of the stimuli we can produce with our pneumatic display. Our overall goal is to integrate this tactile device as a fingertip display mounted on the armature of a force-feedback device, such as the PHANToM [8]. A thorough understanding of the nature of the cues we can deliver will be important to ensuring that we can create constructive, rather than destructive, pairings of tactile pneumatic cues and kinesthetic force-feedback cues. Although not a complete account, the studies we report provide valuable insights as to how this might be done. We have concrete data regarding localization, the two-point threshold and the ability of participants to discriminate the size (or intensity) of presented stimuli. We also have a better understanding of the limitations of our display hardware: we know we cannot render high frequency information. To demonstrate how this psychophysical knowledge can be applied to a practical scenario, in the next section we describe a simple study investigating the utility of combining pneumatic tactile and force-feedback cues to a simple object manipulation task in a virtual environment.

4 APPLIED STUDY – VR BUTTON

4.1 Introduction

This initial investigation focused on the use of haptically rendered VR buttons. These take the form of 3D virtual objects that can (through the medium of a force-feedback device) be pushed, and respond to this action by moving in a manner similar to that of a button in the real world. After an initial resistance is overcome, the button moves back freely for a short distance before coming to a hard stop. This dynamic behavior creates a click-like sensation, a potentially valuable piece of feedback. While it is possible to produce such buttons relatively easily using commercially available software toolkits [e.g. 12], they are not a common feature of virtual environments. Raymaekers & Coninx [11] report one reason for this. They describe a study of a selection task in which participants activated VR buttons by either pushing them (as detailed above), or by using a dedicated physical switch (like a mouse button) positioned on the tip of the end effector of a force-feedback device. This latter condition led to a 50 percent reduction in task completion times, strongly suggesting that requiring VR buttons be pushed in order to be activated is a far from optimal interaction technique.

Despite this negative evidence, creating virtual buttons that behave like real ones remains a relatively desirable goal. If the intention is to immerse users in an environment, or realistically display some complex virtual object to them, then imbuing objects with real world dynamics is a requirement. Although the use of a dedicated hardware button may be more efficient, it is far from a physically realistic approach, and seems likely to destroy any illusions of the realism of the environment. Furthermore in the kind of situations we are considering, in which a user interacts by placing their entire finger in an enclosed haptic display, it is unclear where a physical switch could be mounted.

We suggest that combining pneumatic tactile and force feedback cues relating to button pushing may increase user performance. As inspiration for this, we cite Bicchi *et al.*'s [2] work on tactile flow. One of the fundamental demonstrations of this concept involves improving the accuracy of a feeling of contact by displaying the same net force to a greater or lesser area of skin. By appropriately adjusting the amount of skin being subjected to a stimulus, Bicchi and colleagues were able to create a more realistic percept. Similarly, we may be able to create a more realistic, and potentially easier to use, version of a haptic button by using our pneumatic display to deliver appropriately changing stimuli to the surface of the skin as the button moves through the process of being pressed.

4.2 Experimental Task and Measures

The experimental task in this study was the selection of cube-shaped buttons in a VR environment through the mechanism of moving to the button's surfaces, and pushing against them. To assess performance, we measured the amount of time this took. Two configurations of nine buttons were used, one in which they all faced the user (in which the user activated the buttons by pushing away from his or her body) and the other in which they faced upwards (in which the user was required to push the buttons downwards). In this latter case the clarity of participants' view of the buttons (and therefore possibly their

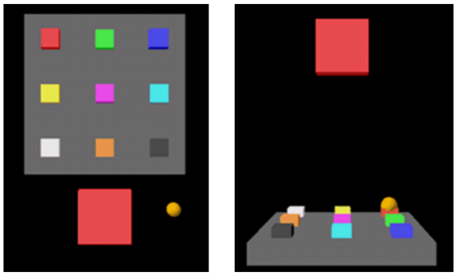

Figure 9. Horizontal and vertical orientations of the virtual environment used in the button study.

task performance) was reduced compared to that in vertical configuration, but the tactile feedback presented by the pneumatic display was more accurate. Subjects pushed downwards (rather than forwards) with their finger, an action which correctly aligns with tactile cues (which are always presented to the base of the finger). Both experimental interfaces are shown in Figure 9. The orientation of the buttons is the only difference in their arrangement; they were sized and spaced identically in a three by three grid pattern. Other than the buttons, and a base object on which they sit, only one other object existed in the scene, positioned at the bottom of the screen in the case of the forward-facing buttons, and the top in the case of the upwards facing buttons. Touching this marker object was used to signify the start of a new trial.

Each of the nine buttons had a different colour (the three primary colors, the three combinations of pairs of primaries, and black, white and orange) and each trial consisted of the subject touching the marker object, at which point it changed from a neutral grey to the colour of one of the buttons. The participant's task was to select the appropriately coloured button as rapidly as possible. When the correct button was selected, both it and the marker object became a neutral grey colour until the start of the next trial. Task completion times were measured from first contact with the marker object to successful button selection.

4.3 Hardware

This experiment was conducted using a modified version of our pneumatic array, and a PHANToM force feedback device [8]. The modifications to the array were simply to add a short finger like protrusion to its tip, which was designed to fit snugly inside the standard thimble attachment of the PHANToM. Thus we were able to augment the normal thimble mode of interaction with the PHANToM (which involves placing a finger inside the thimble) with our pneumatic array with little effort. Problems that might emerge with this ad-hoc solution are that the additional weight it adds to the tip of the PHANToM is unbalanced, and the fact that the user's finger ends up positioned approximately 2 cm in front of the device's fulcrum, the point at which it is designed to accurately apply the forces it generates. These issues are unresolved at the current time, but we believe are sufficiently minor that a meaningful empirical investigation can be undertaken with this hardware configuration. We regard this device as a proof of concept prototype, and suggest that considerable additional development would be required to create a robust, reliable and generally applicable device with the features our preliminary system possesses. The combination of modified pneumatic and PHANToM force feedback device is shown in Figure 10.

4.4 Stimulus design

The movement (and kinesthetic feel) of the virtual buttons used in this study was generated by a simple model consisting of a pair of virtual springs. The first of these was a relatively strong spring active during the initial portion of the button's

Figure 10. Pneumatic array housing for easy attachment to PHANToM force feedback device.

movement, the second a weak spring active during the latter portion. Beyond the range of both these springs, the button's motion simply stopped. This profile results in the following click-like behavior. An initial effort is required to push a button, which (after moving a short distance) drops away rapidly until it reaches a hard stop. This is characteristic of other virtual buttons reported in the literature [11].

The pneumatic tactile feedback we designed to accompany this force-feedback based button interaction was drawn from our previous finding that we can effectively vary the perceived intensity of a pneumatic tactile cue by varying the size of the skin area to which it is presented. It involved altering this factor according to the distance a participant had pushed a button. For the first third of the button's travel we activated the central air jet alone, for the second third the central 9 air jets, and for the final third, all 25 air jets. This resulted in a simple mapping between perceived stimulus intensity and the distance the button had traveled from its initial position; as the distance moved increased, so did the perceived intensity of the tactile cue representing its motion.

4.5 Participants and Experimental Design

All participants from Group B (from the psychophysical studies) completed this study. Two additional participants were also included, bringing the total number of participants to 10. They were both female, one 23, the other 19. The study consisted of a 4 blocks of 90 trials. Each block contained 10 trials for each button, delivered in a random order. The blocks varied according to the orientation of the buttons and presence or absence of pneumatic cues. The blocks containing the forward-facing buttons were always presented before those containing the upwards facing buttons. However, regarding the pneumatic cues, the study followed a balanced repeated measures design. Five subjects experienced a force-feedback condition followed by a force feedback plus pneumatic condition, and the other 5 experienced trial blocks in the opposite order. Prior to each block, participants completed an identically structured practice session consisting of 45 trials.

4.6 Results and Discussion

The results of this study are shown in Figure 11. A 2 by 2 ANOVA was used to analyze the results, and revealed significant effects of both display orientation ($F(9, 1) = 10.68$, $p<0.01$) and the use of pneumatic cues ($F(9, 1) = 6.75$,

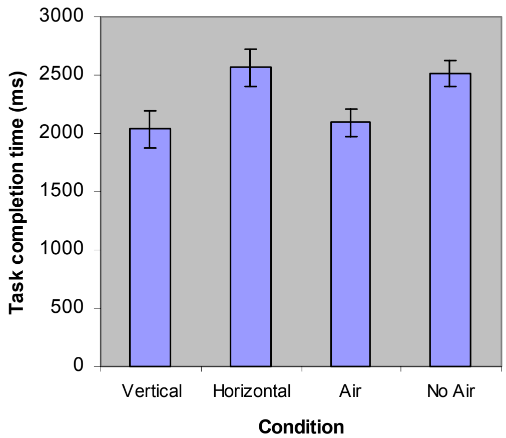

Figure 11. Task completion times in button study.

$p<0.01$). No interaction between these two factors was found. *Post-hoc* t-tests bore out these results for both variables (at $p<0.01$ for display orientation and $p<0.05$ for presence of pneumatic cues).

This study provides simple, but compelling evidence that the combination of proprioceptive, force-feedback cues and cutaneous pneumatic ones can be constructive and yield concrete performance improvements. The addition of tactile cues decreased task completion times by 20% in a targeting task. Considering the simplicity of the experimental task and the frequency with which such fundamental operations are performed during computer use we believe that this observed improvement would translate into significant quantitative and qualitative benefits in more complex scenarios. Indeed, an obvious next step for this work is to perform a comprehensive study around this concept using a complex task and a range of both objective and subjective measures.

Interestingly, the 20% performance improvement was maintained irrespective of the orientation of the buttons; apparently subjects gained the benefits provided by the additional feedback regardless of whether it was a applied to an appropriate portion of the finger or not. Given this apparent spatial robustness and the simplicity of pneumatic cues we presented, an interesting future investigation would be to examine the effect of even simpler, purely binary cues. If performance improvements can be achieved in tasks when contact with virtual objects is presented cutaneously through the activation of a single tactile stimulator, then the kind of improvements we describe can be attained with considerably simpler display hardware.

4 CONCLUSIONS AND FUTURE WORK

In this paper we described the motivations and design of a system that enables the combination of force feedback and pneumatic tactile cues. Through a series of psychophysical studies, we developed a basic understanding of the perceptual characteristics of pneumatic stimuli. We gained insights into localization rates, the two-point threshold, and how intensity is perceived. We also discovered that the limits of tactual temporal perception considerably exceed that of our display hardware. Taking this information on board, we describe a final study which combines proproiceptive and pneumatic cues, and demonstrates that this can lead to performance improvements, and potentially improve the usability of virtual environments.

Development of a better integrated and more robust hardware platform is an imperative for this research; our current device is no more than a proof of concept prototype. Additional future work will involve extending the idea of combining pneumatic tactile and kinesthetic cues to more challenging scenarios. Possibilities we are considering include further fundamental studies using more rigorous procedures and more sensitive metrics and assessment protocols and the enhancement of more complex virtual environments, either featuring larger and more diverse control elements, or those which attempt to achieve greater levels of realism, for instance in the entertainment industry. We believe that the scope of this technology extends to both gaming and narrative entertainment systems, potentially enriching many kinds of digital content and eventually reaching away from the desktop and into the living room.

ACKNOWLEDGEMENTS

This research was supported by the Republic of Korea's Ministry of Information and Communication (MIC) through the Realistic Broadcasting IT Research Center (RBRC) at the Gwangju Institute of Science and Technology (GIST).

REFERENCES

[1] Amemiya, K. & Tanaka, Y. Portable Tactile Display Using Air Jet. In the *International Conference on Artificial Reality and Telexistence*, 1999.

[2] Bicchi, A., Scilingo, E.P. & De Rossi, D. Haptic Discrimination of Softness in Teleoperation: The Role of the Contact Area Spread Rate. In the *IEEE Transactions on Robotics and Automation*, 16, 2000.

[3] Borst, C.W. & C.D. Cavanaugh, *Haptic Controller Deisgn and Palm Sized Vibrotactile Array*, Tech Report, Center for Advanced Computer Studies, University of Louisiana at Lafayette, 2004.

[4] Chodack, J. & Spampinato P. Spacesuit glove thermal micrometeoroid garment protection versus human factors design parameters. In *Proceedings of SAE Conference on Space Station and Advanced EVA*. Society of Automotive Engineers, 1991.

[5] Debus, T., Becker, T., Dupont, P., Jang, T., & Howe, R, Multichannel vibrotactile display for sensory substitution during teleoperation. In *SPIE International Symposium on Intelligent Systems and Advanced Manufacturing*, Newton, MA, Oct 2001.

[6] Johnson, K.O. & Phillips, J.R. Tactile spatial resolution - I. Two point discrimination, gap detection, grating resolution and letter recognition. *Journal of Neurophysiology*, 46(6):1177-1191, 1981.

[7] Loomis, J.M. An investigation of tactile hyperacuity. *Sensory Processes*, 3:289-302, 1979.

[8] Massie, T.H., *Initial haptic explorations with the PHANToM: Virtual touch through pointer interaction*, Masters Thesis, Massachusetts Institute of Technology, Boston, MA, 1996.

[9] Murayama, J., Bougrila, L., Luo, Y.L., Akahane, K., Hasegawa, S., Hirsbrunner, B. & Sato, M. SPIDAR G&G: A New Two-Handed Haptic Interface for Bimanual VR Interaction. In *Proceedings of EuroHaptics 2004*, Munich, Germany, 2004.

[10] Oakley, I, Kim, Y., Lee, J. & Ryu, J. Determining the Feasibility of Forearm Mounted Vibrotactile Displays. In *Proceedings of the Haptics Symposium 2006*, Arlington, VA, 2006.

[11] Raymaekers, C. & Coninx, K. Menu Interactions in a Desktop Haptic Environment. In *Proceedings of EuroHaptics 2001*, Birmingham, UK, 2001.

[12] Reachin Technologies AB, http://www.reachin.se/, January 2006.

[13] Sekuler, R. & R. Blake, *Perception*. Mcgraw-Hill College, 2002.

[14] Tan, H.Z., Srinivasan, M.A., Eberman, B., & Cheng, B. Human factors for the design of force-reflecting haptic interfaces. In *Proceedings of ASME Dynamic Systems and Control Division*, Chicago, IL, 1994.

[15] Van der Linde R.Q., Lammertse P., Frederiksen E. & Ruiter B. The HapticMaster, a new high-performance haptic interface. In *Proceedings of EuroHaptics 2002*, Edinburgh, UK, 2002

[16] Verrillo, R.T., Effect of Contactor Area on the Vibrotactile Threshold, *Journal of the Acostical Society of America*, 35(12):1962-1996, 1963.

[17] Wagner, C.R., Clatz, O., Perrin, D.P., Feller, R.L., Delingette, H., Ayache, N. & Howe, R.D., Integrating Tactile and Force Feedback with Finite Element Models. In *IEEE International Conference on Robotics and Automation*, Barcelona, Spain, 2005.

[18] Webster, R., Haluck, R.S., Zoppetti, G., Benson, A., Boyd, J., Charles, N., Reeser, J. & Sampson, S. A Haptic Surgical Simulator for Laparoscopic Cholecystectomy Using Real-Time Deformable Organs. In *Proceedings of the IASTED International Conference of Biomedical Engineering*, Salzburg, Austria, 2003.