WindyWall: Exploring Creative Wind Simulations

David Tolley[1] Thi Ngoc Tram Nguyen[1] Anthony Tang[2]
Nimesha Ranasinghe[1,3] Kensaku Kawauchi[4] Ching-Chiuan Yen[1,5]
[1]Keio-NUS CUTE Centre
[4]Smart Systems Institute
[5]Division of Industrial Design
National University of Singapore
[2]Department of Computer Science
University of Calgary
[3]School of Computing and
Information Science
University of Maine

ABSTRACT

Wind simulations are typically one-off implementations for specific applications. We introduce WindyWall, a platform for creative design and exploration of wind simulations. WindyWall is a three-panel 90-fan array that encapsulates users with 270° of wind coverage. We describe the design and implementation of the array panels, discussing how the panels can be re-arranged, where various wind simulations can be realized as simple effects. To understand how people perceive “wind” generated from WindyWall, we conducted a pilot study of wind magnitude perception using different wind activation patterns from WindyWall. Our findings suggest that: horizontal wind activations are perceived more readily than vertical ones, and that people’s perceptions of wind are highly variable—most individuals will rate airflow differently in subsequent exposures. Based on our findings, we discuss the importance of developing a method for characterizing wind simulations, and provide design directions for others using fan arrays to simulate wind.

Author Keywords

Multimodal Interaction; Novel Actuators / Displays; Tactile / Haptic Interaction

CSS Concepts

• Human-centered computing~Haptic devices

INTRODUCTION

Virtual Reality (VR), simulations have long focused on sight and sound as primary modalities for conveying immersion, or presence in a remote environment [38]. Recent work has begun to explore other sensory modalities, including haptics, temperature, olfactory, and taste (e.g. [32]). These modalities have been demonstrated to increase users’ sense of presence in virtual spaces, with the impact of multiple combined modalities having a strong gestalt effect on immersion beyond each modality on its own. The present work lays down groundwork for VR simulations intended for relaxation, learning and therapy [10, 14, 30, 7], since effecting a strong sense of presence is crucial to the effectiveness of these simulations.

Figure 1. WindyWall is a 3-panel array for simulating wind.

Our focus is on designing wind simulations, where we are generating airflow on and across skin, and so we are interested in flexible ways to explore airflow simulations and understanding how people perceive these simulations. For instance, we are interested in being able to simulate scenarios such as a small puff of air, a gentle rolling breeze, a sweeping wind and an object rushing past the user. While knowing how to design wind simulations is useful for such VR experiences, this knowledge will also prove useful for a number of other applications, including ambient awareness [36, 13, 37], notifications [20, 17, 24, 27], and even subtle cuing of attention (e.g. [23]).

The problem is that we do not have a reproducible method for simulating wind, or a sufficiently nuanced understanding of the psychophysical experience of airflow for designing effective simulations. Wind as a modality is still largely understudied by the research community: most work has been limited to proof-of-concept systems for specific wind simulations and installations [1, 5, 8, 12, 18, 19, 21, 24, 25, 27, 35, 32], and as wind displays for feedback [17, 20, 29]. While these are fascinating, it is difficult for others to build on and creatively explore these ideas without the specific hardware in hand. Further, little work has explored how people experience airflow: for instance, how people experience the magnitude of airflow, whether the way airflow is directed at them affects this, or how people determine the source direction of airflow (and how accurate they are). Addressing these questions will help us in designing effective wind simulations that produce the sensation that designers are trying to achieve.

TEI '19, March 17–20, 2019, Tempe, AZ, USA

ACM 978-1-4503-6196-5/19/03…$15.00
https://doi.org/10.1145/3294109.3295624

We take a three-pronged approach to address these problems: first, we introduce WindyWall, a general platform for exploring creative wind simulations; second, we develop an initial approach for characterizing specific elements of a wind simulation, giving others a way to parameterize and compare different wind simulations, and finally, we design and conduct a pilot experiment to understand how people perceive airflow as a function of how we manipulate parameters of WindyWall. WindyWall is a 90-fan array divided between three separate panels, where panels can be flexibly re-arranged, and individual fans can be actuated independently depending on the demands of a particular simulation. For instance, we can simulate a point-source gust of wind, or a column of wind, or modulate fan actuation timing to simulate more dynamic wind patterns.

Our pilot study informs future work with wind simulations by assessing users' perception of airflow magnitude. Our findings show that users are more sensitive to horizontal activation patterns of simulated wind (compared to vertical activation patterns), but that there are wide variations in how people perceive magnitude in airflow. These findings suggest that wind simulations need to be carefully designed and tested, and that our naïve expectations of how airflow is experienced may not be borne out in reality.

In this paper, we make the following contributions: first, we contribute the design of WindyWall, a fan array testbed to realize wind simulations; second, we outline requirements for characterizing wind simulations, and finally, through a pilot study, we provide new information to designers of wind simulations regarding how people perceive airflow as a function of different wind display parameters.

RELATED WORK

Simulating Wind for VR and Other Applications

Many researchers have used airflow as a modality to enhance the feeling of presence in virtual environments; however, the field has mainly produced a series of one-off system designs that are difficult to reproduce, each with its own idiosyncrasies. Some wind displays are embedded in environments [25, 36, 8], while others are embedded in everyday contexts, such as mounted to an existing display monitor [27, 24], or head-mounted [5, 17]. Collectively, these systems do not clearly build from one another, either in terms of design lessons, or technical approach; rather, each seems like a standalone, proof-of-concept design. While a comprehensive review is beyond the scope of this work, we outline exemplars of each of these approaches.

Full Body Displays. Some systems take a person-scale or room-scale approach to designing the airflow display. Moon & Kim [25] designed a 1m×1m×2m frame where 20 fans were affixed at three different heights. The authors found that fans blowing at the user's torso and head were most appropriate, whereas fans blowing at users' legs felt as though wind was blowing upward. In a study where participants viewed videos playing externally to the frame, the authors found that participants felt higher levels of immersion and presence in the VR simulation when the wind display was active. At an even larger scale, Verlinden et al. [36] report on a room-scale wind simulation, where fans are affixed to a frame hanging above a CAVE-like environment. The entire room was used for simulating sailing, where wind is an important decision-making modality for sailors. Participants in a study reported feeling more immersed in the simulation with the wind display. Such context-specific approaches have been demonstrated to improve performance. For instance, Deligiannidis & Jacob [8] show that wind and tactile feedback for a scooter simulation improve the speed through which participants complete a virtual obstacle course.

Augmenting the Environment. Some research augments conventional environments with fans to explore the design space of wind displays. For instance, Minakuchi & Nakamura [24] study mounting fans around a conventional computer monitor, and on the back of an office chair. Similarly, Mowafi et al. [27] assembled an array of six fans and nozzles along the top of a monitor to explore the extent to which people could differentiate between different patterns of activation. These kinds of approaches are more tractable explorations of wind displays, though not as common as head-mounted approaches.

Head-Mounted Displays. Head-mounted airflow displays align with the increasingly common approach to use head-mounted VR goggles. Ambiotherm, for instance, uses a pair of articulated fans mounted below the goggles to simulate wind [32]. In a study of the impact of multimodal stimulation (including temperature, sound, etc.), the authors demonstrated that using the fan to stimulate wind increased feelings of immersion. Other researchers use multiple actuators for more spatially localized airflow simulation. Cardin et al. [5] use multiple fans mounted on a specialized head-mounted frame, where eight different fans could be actuated at different times to simulate wind from different directions. Kojima et al. [17] take a similar, helmet-based approach, where the airflow is generated by using audio speakers to push air through small tubes at users' ears.

Understanding Perception of Airflow to Design Effective Wind Displays

To effectively design with wind as a modality, we need to consider not only the technical aspects of building these airflow displays, but also gain a clear understanding of how airflow is perceived by people. Unfortunately, our current understanding of this is a patchwork of experimental results from one-off systems.

Magnitude Estimation. Agdas et al. [1] assessed peoples' ability to discern wind speed with a large turbine-style fan in a trailer, where airflow was simulated at rates of 10-70 mph. While the work was conducted in the context of public policy (i.e. are people effective in assessing the wind speed, and the risk of corporeal or property damage), the results are informative for design. Of particular interest, the authors observed a Weber-Fechner effect, where people's assessment of wind speed are non-linear to increases in actual wind speed. It would be interesting to understand how this relationship plays out at more appropriate wind speeds reasonable for VR simulations.

Detection Thresholds. Lee & Lee [20] report on a series of experiments to understand detection thresholds for airflow positioned and aimed at different parts of one's body. For instance, at what intensity does the airflow need to be detected (i.e. noticed) by a participant when the airflow is directed at different body parts (e.g. cheek, wrist, ankle, foot, etc.). Beyond intensity, the authors also explored duration threshold (i.e. how long does the user need to be exposed to airflow to detect airflow), and the distance threshold (i.e. how close does the device need to be for detection).

Directionality. Nakano et al. [29] present a study exploring how well people can differentiate between different fans, in angular terms, around their torso. Their study relied on a seven-fan array where fans were arranged 20 degrees apart from one another in a semi-circle with a radius of 80cm. This fan array was arranged around a participant's head (where the participant's head was fixed to a mount). The authors used a just-noticeable difference (JND) study design, where participants indicated whether two fans were the same or different fans. The authors report two JND values of 0° and 30°, depending on whether the airflow came from head-on, or the side. However, the authors also report on a small misalignment of one of the fans, which could have caused wide variation in their results. Notably, participants used different strategies to determine the source of the airflow; while some relied on sensations on their noses, others relied on sensations on their cheeks.

Synthesis. We have already seen a wide range of airflow displays designed for different scales; however, we do not yet have a good way of characterizing and parameterizing the wind displays (e.g. what are their characteristics of each simulation), nor do we yet have a coherent articulation of the psychophysical experience of airflow, particularly as it relates to design parameters that we can control. We do not know, for instance, how people perceive the magnitude of airflow when the contact surface area is varied. To address these issues and shortcomings, we designed WindyWall, a wind array that allows us to control various factors and understand how airflow is perceived by users.

WINDYWALL PLATFORM

As illustrated in Figures 1 and 2, WindyWall is a hardware and software platform comprising of three independent panels of fans. Each panel comprises of 30 individual fans arranged in a 5×6 array (width × height), where each fan can be individually actuated. Each panel can be individually re-arranged or placed around a user in a 270° arrangement.

Design Goals

We are motivated by an increasing interest in developing immersive VR simulations that incorporate wind. As such, we were initially interested in designing a platform that could simulate a range of wind scenarios, including a small puff of air, a gentle rolling breeze, a sweeping wind (e.g. a pedestal fan) or an object rushing past. Beyond this, our overriding goal in constructing the WindyWall platform was to design a platform to enable designers to creatively explore different possibilities within wind simulation. A reconfigurable platform would allow the designer to explore different types of setups, while the software interface would allow for semi high-level programming of the hardware platform.

Based on these ideas and our exploration of prior work, we developed four design goals for the WindyWall platform:

DG1: The platform needs to be designed for a single-person simulation for wind.

DG2: The platform must enable individually actuated sources of airflow, where each can be independently set to a specific power level.

DG3: The platform must be physically reconfigurable to allow for creative explorations of wind simulation.

DG4: Finally, the software platform ought to provide an easy mechanism to control groups of individually activated fans, where effects can be designed spatially and temporally.

Design and Implementation

Our primary interest in this iteration is the design of a single-person simulation platform (DG1). While larger setups exist for different kinds of simulations (e.g. for groups, or for individuals), we were interested in designing for setups where an individual is immersed in a VR simulation, facilitated primarily through head-mounted stereo displays, rather than wall-projected CAVE systems. Such a setup is more portable, making it more suitable for traveling demonstrations as necessary. Furthermore, the reduced physical footprint would make it easier to replicate in multiple sites when designing multi-user experiences.

The central design tension we needed to resolve was how reconfigurable the platform ought to be (DG2, DG3, DG4) with regard to the ease of managing and designing a simulation. For instance, a fully-reconfigurable setup could involve individually positioned fans, each on a flexible arm clamp, but consistently recreating this setup would be time-consuming and onerous. In our process, we realized that if we were simply designing the platform to be experienced by a single individual at a time, we only needed to focus on the airflow experiences of the individual in that context—thus, it was only necessary to focus on fans whose airflow would touch the individual directly. To resolve the design tension, we settled on a partially reconfigurable set of panels (DG3), where each panel would contain a standard set of fans. These fans would all be pointed toward the user of the system, but individually not be reconfigurable within the panel itself. On the other hand, the software platform would enable the designer to access and control each of the individual fans (DG2), while a thin software layer could programmatically give the designer access to rows, columns, or arbitrarily assigned sets of fans as a group (DG4).

The final design has three panels providing 270° of coverage around a single user. Since the panels are easily reproducible and reconfigurable, they also allow other kinds of arrangements. Some possibilities include: (a) fully

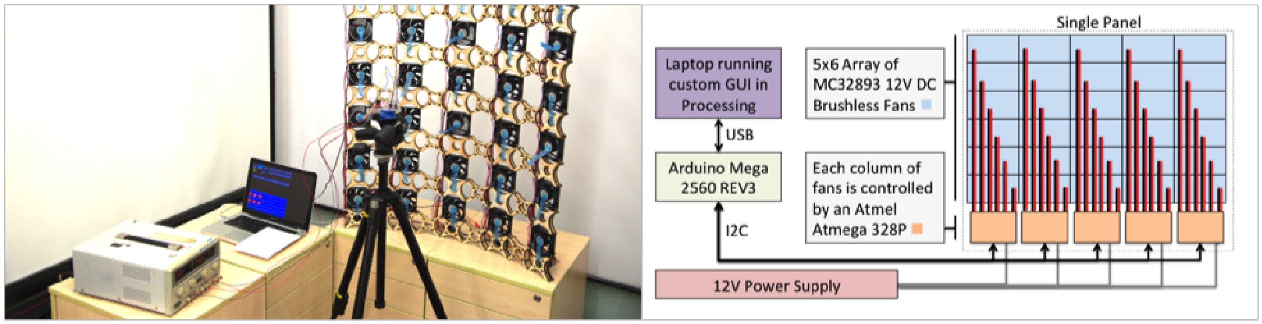

Figure 2. The testing configuration for a single WindyWall panel (left). The schematic of the WindyWall (right).

surrounding a user with panels; (b) a "hallway" of panels that allow a user to move through the space; (c) a "room" of panels, and (d) additional panels suspended from the ceiling to provide a "top down" sense of airflow as required.

Each panel comprises of 30 fans connected horizontally and vertically by wooden joints. Joints are 5mm thick pieces of laser-cut 5-ply plywood, connected to the corners of each fan. Angled at ~18° each, the horizontal joints connect each column of fans in an arc shape to form a curved wall, which covers an overall angle of 90°. Within the array, wooden joints create a horizontal spacing of 9.5cm and a vertical spacing of 8.5cm between each fan. When connected in this manner, each panel has a total height of 104.5cm and a curved width of 80cm. Each of the 90 fans is an MC32893 12V DC brushless fan, manufactured by Multicomp [28]. Individually, these fans are 80mm x 80mm x 25mm in size, have a rated power consumption of 1.7W and a rated airflow of 41CFM. Each column of 6 fans is connected to an Atmel ATmega328P [3] microprocessor and, in turn, each of these chips are connected to an Arduino Mega 2560 REV3 microcontroller [4], which acts as a "master" chip. Utilizing I2C protocol, both singular and multiple fans can be controlled through a custom GUI that has been created in Processing [31]. When a user wants to actuate specific fans, the Arduino Mega sends commands to the appropriate ATmega328P chips that are linked to the address of those target fans. Once each ATmega328P receives a command it will analyze it and execute any changes to fan-rotation speeds that are specified by the command.

Designed Wind Effects

WindyWall is over-engineered to support our interest in designing wind effects. The platform was designed to give designers a wide range of capabilities for flexibly and creatively exploring the design of wind simulations. We expect, for instance, that the granularity of the individual fans to be far higher than what can be discerned by most human users; thus, a practical system for deployment might make use of far fewer actual fans. Furthermore, other kinds of wind simulations (i.e. for more than one person) would need different kinds of implementations. Nevertheless, to illustrate our approach, we designed several wind effects that could be conceived of as small scenario-based wind simulations:

- **Puff of Air**: A small number of closely located fans are simultaneously actuated at high fan speed for a short period of time. This can be used to simulate a short human breathe near a user's ear or an abrupt rush of air as a door is slammed shut near the user.
- **Rolling Breeze**: The fan speed of multiple fans is constantly modulated within a set range over time. This can be used to simulate continuous fluctuations of wind speed that are commonly experienced in outdoor environments. Here, this technique can be combined with broader or narrower configurations of fans to provide more or less ambient wind coverage as required (e.g. when simulating a user transitioning between two environments, such as walking in and out of a house).
- **Sweeping Wind Movement**: Adjacent fans are continuously actuated one after another to create a wind simulation that can sweep between any two points covered by the fans. This technique can be used to simulate the movement and orientation of the user or surrounding objects within a virtual environment. For example, a virtual rotating desk fan may be simulated by moderately-paced wind sweeps oscillating in the transverse plane. Alternatively, a sweeping motion may also be used to imply a user's movement past a virtual wind source (e.g. a user walking past an air vent).
- **Objects Rushing Past the User**: By combing the Sweeping Wind Movement and Puff of Air techniques, it is also possible to create more fast-paced sweeping wind simulations. These may be applied to scenarios where objects rush past a user in close proximity (e.g. a user dodging virtual projectiles that are being thrown at them).

TOWARDS CHARACTERIZING A WIND SIMULATION

To compare various implementations of a wind simulations (e.g. each of our wind effects), we need a standard way of

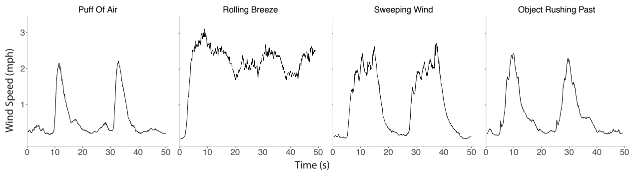

Figure 3. Our characterization technique captures airflow magnitude over time, and visually allows us to distinguish between different airflow simulations. Left to right: Puff of Air, Rolling Breeze, Sweeping Wind and Object Rushing Past.

characterizing how the airflow in the simulation is generated. Such a standardized method should allow researchers to describe and compare between setups and wind simulations (e.g. configurations, or effects) in an objective way. The method should emulate the experience of wind from the perspective of a human participant, measuring airflow, accounting for the ramp-up time of the airflow generator (e.g. fan, or other device), the sustained airflow rate (e.g. L/s), a way of describing the directionality of airflow (allowing for multiple directions of airflow), and how this airflow changes in relation to time (e.g. for an effect such as a gust of wind).

We developed a characterization strategy that satisfies the first two of these requirements.

Measurement Device: We use a thermal anemometer as our measurement device. The thermal anemometer (Rev P Wind Sensor [2]) relies on a hot-wire wind sensing mechanism, giving us appropriate resolution and consistency for the lower wind speeds generated by WindyWall (0-3 mph). We constructed a redundant measurement device using a pair of sensors, where each sensor samples wind speeds at 10Hz. In our setup (Figure 2, left), this device was affixed atop a tripod, and positioned to approximate the location of a seated user's nose (~60cm from the center fans).

Collecting Data: We collect wind speed data (mph) in an enclosed, temperature-controlled room (22°C) with minimal air circulation. In a typical situation, we will run the simulation multiple times (e.g. three times), collapsing the data across these multiple runs to get a more representative result, giving a clearer picture of the overall nature of a simulation apart from individual one-off sensor variations from a given measurement session.

Visualizing Data: Plotting the collapsed data as a time series plot shows the temporal windspeed variations, illustrating how each of the simulations are characteristically different.

To illustrate our approach, Figure 3 visualizes the characterization charts for the four wind effects described in the previous section. These show that Puff of Air is much briefer than Object Rushing Past, both of which are much shorter than Sweeping Wind and Rolling Breeze. Further, the peaks have a higher kurtosis than with the Rolling Breeze, which means the peak blowing-speed is briefer in Puff of Air than in Rolling Breeze. Finally, Sweeping Wind takes longer to reach its peak blowing-speed compared to Puff of Air. Thus, this simple visualization shows us ramp-up speeds, holds, sustained airflow rates, pauses, and ramp-downs in each airflow simulation. Even here, this technique provides designers a simple way to compare between different creative simulations with WindyWall.

While this strategy provides the first steps toward a wind characterization method, there is still considerable room for improvement. For instance, our technique does not yet account for airflow directionality, nor provide a broader "surface" for capturing airflow. Puff-of-air and object-rushing-past, for instance, look similar, even though an object rushing past sweeps by alongside one's body. This distinction is not captured in the visualization above, even though they are subjectively fairly different experiences. Currently, the measurement device only measures airflow from a single point, but we experience airflow along our entire bodies, thus a device that provides a larger surface area for air to impact it would be useful. Critically, future approaches need to emulate the experience of a human participant, which our strategy attempts to capture. This characterization technique should open the door to parameterizing wind simulations, allowing us to compare between different simulation platforms and effects.

PILOT STUDY OF AIRFLOW MAGNITUDE PERCEPTION

Beyond objective characterizations of WindyWall airflow simulations, we are also interested in the subjective experience of airflow. To date, there is still very little work in this space, and since we are ultimately interested in designing wind simulations as part of immersive multisensory experiences, there were several questions that we wanted to address: for instance, is there a clear relationship between actual objective airflow rate and perceived airflow magnitude by users? Is this affected by how the airflow is directed at the user? To what extent, given repeated exposures, are individuals consistent in their assessments of the same airflow simulation? Across a group of users, are individuals consistent in their assessments of the same airflows? Answers to these basic questions are critical to designing simulations that can be dependably experienced as intended by the simulation's designer.

To address these questions, we designed a study where users' perceptions of airflow magnitude are compared to objective measures across several simple wind simulations. Our experiment asked participants to provide ratings of airflow—"How much wind do you feel?" (i.e. perceived magnitude of airflow) on a self-reported scale. We varied two parameters with WindyWall across airflow simulations:

- **Fan Rotation Speed**: High and Low (40% and 80% power, respectively, determined by early pilot studies for reliable and distinguishable fan speeds).
- **Size and Shape of Airflow Source**: Changing the number of fans, and the shape of the airflow simulation.

While studies have been conducted to explore this phenomenon given very high wind speeds (e.g. 10mph-60mph) [1], our focus was centered on investigating more tractable levels of airflow generation for the purpose of multisensory immersion (i.e. wind speeds of 0 – 3mph). For us, we wanted to compare wind blown at head level to match approaches from prior work (e.g. [29, 32, 24, 27]). Beyond this, we were also interested in exploring whether different spatial arrangements of active fans affected the perceived magnitude of the airflow to extend the arrangements explored from prior work (e.g. [25, 36]).

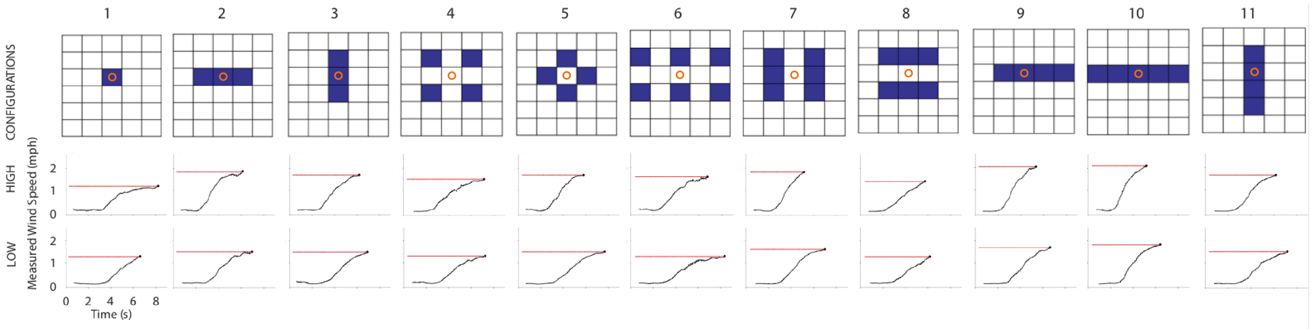

Figure 4. Tested fan configurations at low and high power. The first row illustrates configuration on a single panel (dark = fan on; white = fan off; circle = central fan). Each chart illustrates the ramp up function for the fan array to reach the sustained wind speed (red line).

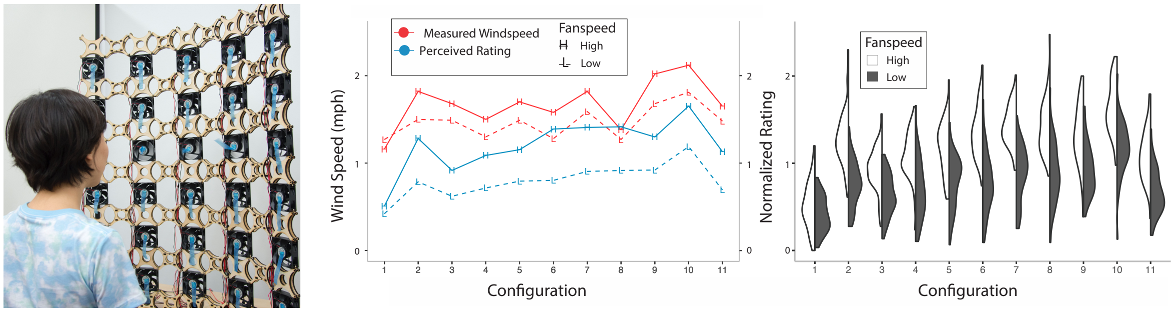

Figure 5. The magnitude estimation experimental setup (left). Average perceived magnitude plotted against measured wind speed for each configuration and fan speed (center). Distribution of ratings for each configuration and fan speed (right).

Objective Data Collection. To simplify the execution of our study, we used a single WindyWall panel, and very simple simulations (i.e. on/off). We explored eleven different fan configurations (illustrated in Figure 4), where the central fan was aimed directly at the sensor (or participant's head). Early pilots determined that the top and bottom two rows of fans could not be reliably felt if a participant was sitting. Thus, the fan configurations vary the number of fans (i.e. one fan to six fans), and also in different variations (i.e. horizontal, vertical, narrow, wide).

We collected wind speed data (mph) for the 11 fan configurations three times, each running at two different power levels (40% and 80%) that represented low and high fan speeds. This objective data is visualized in Figure 4.

Pilot Study Design and Procedure. Using a repeated-measures design, participants were asked to rate the same 11 configurations × 2 speeds a total of three times over the course of two weeks. The presentation order of the 22 simulations was randomized for each participant, with each configuration exposure time lasting for 20 seconds.

We recruited 16 participants (7 women) from our university, aged 19–33 (M=25, SD=3.39). Participants wore short-sleeve T-shirts to ensure consistent skin exposure. If participants normally wore glasses, they were asked to remove them for this study. As shown in Figure 5 (left), participants were seated facing the WindyWall panel at a distance of ~60cm. The height of the chair was adjusted so that the participant's nose was at the same height as the center fan of the third row. Participant were asked to close their eyes, and listened to white noise to cover fan noises.

Each round of exposure comprised of three events: (1) *Fans Activated*: Fans are actuated for the airflow simulation in this round, and given time to stabilize (5 seconds); (2) *Participant Provides Rating*: A bell sound plays, and participants provide a verbal rating of the perceived wind magnitude (10 seconds); (3) *Fans Deactivated*: Fans are allowed to come to rest (5 seconds). When providing ratings, participants were asked to rate the magnitude of the wind using the first wind simulation as a basis, subsequently providing ratings as ratios compared to their first exposure. In keeping with the *magnitude estimation procedure* outlined by [15], we did not limit the range that the participants could report, instead allowing participants to define their own ranges and to assign any number of their choice. To limit fatigue, we designed the study so that each rating exercise lasted no more than 15 minutes.

Analysis. As participants were free to give their ratings on any scale, we used a normalization technique to bring ratings into a consistent scale across the entire sample. Consistent with the process outlined in [15], raw ratings were normalized in two steps: first, we divided all ratings by their first one to give a consistent unit rating; second, we normalized all ratings such that an individual's mean ratings for 22 configurations are equal to 1.

Pilot Study Results

We present preliminary findings from our pilot study, focusing on: the relationship between actual and perceived airflow intensity, the consistency of an individual's ratings and the consistency of ratings between individuals.

Actual vs. Perceived Airflow Magnitude. Figure 5 (center) shows a moderately positive correlation between participants' perceived airflow magnitude and objective measurements for each simulation (Pearson correlation coefficient=0.78). This suggests participants were able to perceive increases in airflow to some extent: wind magnitude ratings for airflows produced by high fan speeds were higher than those at the low fan speed. Participants' ratings of perceived airflow magnitude generally seemed to increase when exposed to a greater number of fans. As shown in Figure 5 (center), the simulations with the largest magnitude ratings all used either 5 or 6 fans.

Simulations with fans arranged on the transverse plane (horizontal around one's face, from one cheek to the other) were perceived to have a higher magnitude than those arranged in the sagittal plane (vertical from the top of the forehead to the chin). For instance, although configurations 2 and 3 both consisted of three fans, paired samples t-tests showed that configuration 2 (three horizontal fans) had significantly higher perceived ratings than configuration 3 (three vertical fans) at both high and low fan speeds, respectively: 2-high (M=1.28, SD=0.363) vs. 3-high (M=0.92, SD=0.252), $p < 0.005$; and 2-low (M=0.78, SD=0.254) vs. 3-low (M=0.62, SD=0.231), $p < 0.005$. Similarly, configurations 9 and 11 both consisted of four fans, and paired samples t-tests showed that configuration 9 (four horizontal fans) had significantly higher perceived ratings than configuration 11 (four vertical fans) at both high and low fan speeds, respectively: 9-high (M=1.298, SD=0.358) vs. 11-high (M=1.13, SD=0.349), $p = 0.014$; and 9-low (M=0.92, SD=0.309) vs. 11-low (M=0.69, SD=0.260), $p < 0.0005$.

Reliability of an Individual's Ratings. Individuals' ratings were either poorly or moderately reliable across their three sessions. Using Intraclass Correlation Coefficient (ICC), we measured each participant's reliability across multiple exposures for each configuration—i.e. intra-rater reliability [16]. Our analysis used 22 configurations, three measures (k=3), a two-way mixed model, absolute-agreement and single measurement as parameters. Across participants, ICC ranged from 0.37 to 0.76 (M=0.60, SD=0.097). Of our 16 participants, 3 participants' reliability would be considered poor (ICC < 0.50), 12 were considered moderate (0.50 < ICC < 0.75), while only one was considered good (0.75 < ICC < 0.90).

Consistency Across Individuals' Ratings. We found that participants did not agree on the perceived magnitude of each simulation. Figure 5 (right) illustrates the distribution of ratings across all participants for each simulation. The high amount of dispersion (i.e. lack of tightness around the central value) for each configuration indicates that agreement between participants is quite low: for any given simulation some participants rate the experience as being windy, while others might rate it as not being that windy. Many of the configurations seemed to provide a similar "magnitude" experience, even though naïvely, we would not expect this. As illustrated in Figure 5 (right), the distributions for many configurations share a lot of overlap, meaning that many different configurations essentially felt "the same" to participants. For instance, we might expect configurations 7 and 8 to feel similar given the similarity of the arrangement of the fans, and this is demonstrated in the distributions. On the other hand, configurations 2 and 6, which have different numbers and arrangement of fans, also appear to be experienced similarly by participants.

Limitations

These findings should be considered only provisionary. First, we did not examine extended periods of repeated exposure (multiple repetitions per session). Second, the overall airflow volume produced by WindyWall may be too low for participants to distinguish various simulations. It may be the case that, for higher airflow volumes, participants may be more effective at distinguishing between different wind stimuli. This suggests that future iterations of WindyWall, or other wind interfaces, ought to be sufficiently higher powered to provide a greater range of airflow volumes that can be distinguished by users.

DISCUSSION

Characterizing Wind Simulations

We reiterate that it will be important to develop effective techniques to characterize wind simulations for reproducibility and comparison across different setups. Our efforts took initial steps toward developing this method, but we think there are still steps to be taken in experience capture, characterization, and visualization.

Higher Fidelity Experience Capture. Our major insight was that designers ought to try to characterize a wind simulation from the perspective of a typical human user. To this end, our capture device was a wind sensor made of two small anemometers placed where a user's face would be positioned when using WindyWall. Yet, a human's body is likely to have many places where it is sensitive to wind. If we assume that sensitivity is mainly in the face or ears [20, 29], then it may be more appropriate to outfit a mannequin bust with several such wind sensors (e.g. at the ears, at the nose, on the cheeks, on the forehead, neck, and so forth). Such an approach would allow us to capture not only a single point of data about airflow, but a mesh-style dataset accounting for airflow variation around a face.

Further, our lived experience suggests that bodies can differentiate between single and multiple sources of wind on the same patch of skin. For instance, if one was to have two separate fans blowing, this would feel different from a single fan blowing on the same patch of skin—not just in terms of airflow (i.e. volume), but also in that it actually feels like the wind is coming from two different sources. Currently, our anemometers only report a single airflow direction and airflow rate. While these are useful points of data, the sensor still remains a far cry from the sensitivity of our skin.

Higher-Dimensional Characterization. To effectively characterize wind simulations, we need to account for more than just airflow direction and volume: we also need to account for variations of the two parameters over time, and at a given moment, note that different places of our body may be experiencing different airflow direction and volume. For instance, consider the experience of a car that passes by quickly from one's left to right. Even if we do not touch this car, we do feel perturbations in the air due to fluid dynamics, and we feel this as a type of wind or airflow. Over time, this experience changes—we first feel it on the left side of our body, and as the car passes, we feel the wind on the right side of our body. The methods we use to characterize airflow need to be able to describe this change over time—not as theoretical constructs as in a fluid dynamics equation, but in terms of the experience as a human perceives it. Similarly, different points of our body experience different rates of airflow, and due to fluid dynamics, each part of the body may be experiencing the airflow from different directions.

Distance from Fans. Finally, how far people are from fans affects their experience of airflow. While in our work, we strictly control this factor (by placing people 1m away from fans), this will not be generally true in most real-life deployments. For instance, in the case of immersive VR experiences, people will be holding controllers such that their arms and hands will be closer than 0.6m from fans. This needs to be considered carefully in future work.

Implications for Immersive Wind Simulations

Our explorations and findings suggest that people are far less capable of distinguishing and describing airflow than we had anticipated, that we need to think carefully about how people experience wind, and that our use of fans for simulation method may not be optimal.

Variability in Sensitivity. Our results suggest that people's sensitivity to airflow is highly variable—at least in terms of magnitude estimation: some are good at detecting subtle variations of airflow, while others are largely unable to detect meaningful variation. This presents challenges to designers of wind simulations, as it is difficult to know how to simulations are going to be experienced by a population of users. This suggests that we need to do more as a research community to explore the psychophysical experience of airflow, or how this modality acts in concert with others for virtual reality simulations.

Passive vs. Active Exploration. Most research in this area has assumed that people discover wind passively: they sit in a fixed location, and have wind blown at them. Yet, in a virtual simulation, a user will stand, move around, and try to explore the space (and thus experience wind) in an active manner. We still do not have enough understanding of how to account for and accommodate these kinds of situations, and thus do not know how to design effectively for it.

Challenges with Fans. Our experiences suggest that while fans are a cheap and easy method for simulating wind, this may not be the best approach. Fans have a "ramp-up" time (illustrated in Figure 4), adding to the latency in a simulation—for us, recorded latencies ranged between 2 and 4 seconds. Fans also produce vortices of wind, which may be undesirable in certain kinds of simulations. Recent efforts in developing haptic experiences (e.g. for feeling virtual objects, notifications, or force feedback) have explored the use of ultrasound and air pumps [22, 33, 34, 11]. It may be possible to use technologies in combination with fans to simulate foreground and background experiences of wind.

CONCLUSIONS AND FUTURE WORK

WindyWall is a reconfigurable fan array that gives designers the power to creatively explore wind simulations. Our design allows the reconfiguration of the physical array, and allows each fan to be controlled individually, or in programmatic groups. We detail the issue of characterizing wind simulations, and its importance to the community as a whole. To understand the human experience of wind, we conducted an initial pilot study exploring the psychophysical experience of wind and airflow magnitude. Our results suggest that while overall trends exist, individual variation is extremely high. Our findings suggest that people may not be able to articulate their experience of wind effectively; instead, perhaps coarse-grained approaches may actually be sufficient for many applications. WindyWall opens up the possibility to explore these applications and ideas.

ACKNOWLEDGEMENTS

This research is supported by the National Research Foundation, Prime Minister's Office, Singapore under its International Research Centres in Singapore Funding Initiative, and the National Science and Engineering Research Council of Canada. We thank Ivan Tan for supporting this work. We also thank participants of our study.

REFERENCES

[1] Duzgun Agdas, Gregory D. Webster, and Forrest J. Masters. 2012. Wind speed perception and risk. *PloS one* 7. no. 11 (2012): e49944.

[2] AS220 Industries. https://moderndevice.com/product/wind-sensor-rev-p/.

[3] Atmel. https://en.wikipedia.org/wiki/ATmega328.

[4] Arduino. https://www.arduino.cc/en/uploads/Main/arduino-mega2560_R3-sch.pdf.

[5] Sylvain Cardin, Daniel Thalmann, and Frederic Vexo. 2007. Head mounted wind. In *Proceedings of the 20th annual conference on Computer Animation and Social Agents (CASA2007)*. no. VRLAB-CONF-2007-136, 101-108.

[6] Tom Carter, Sue Ann Seah, Benjamin Long, Bruce Drinkwater, and Sriram Subramanian. 2013. UltraHaptics: multi-point mid-air haptic feedback for touch surfaces. In *Proceedings of the 26th annual ACM symposium on User interface software and technology*. ACM, 505-514.

[7] Amber Choo, Xin Tong, Diane Gromala, and Ari Hollander. 2014. Virtual reality and mobius floe: cognitive distraction as non-pharmacological analgesic for pain management. In *Games for Health 2014*. Springer, 8-12.

[8] Leonidas Deligiannidis, and Robert JK Jacob. 2006. The vr scooter: Wind and tactile feedback improve user performance. In *IEEE Symposium on 3D User Interfaces (3DUI 2006).* IEEE, 143-150.

[9] Mark H. Draper, Erik S. Viirre, Thomas A. Furness, and Valerie J. Gawron. 2001. Effects of image scale and system time delay on simulator sickness within head-coupled virtual environments. *Human factors* 43, no. 1 (2001): 129-146.

[10] Diane Gromala, Xin Tong, Amber Choo, Mehdi Karamnejad, and Chris D. Shaw. 2015. The virtual meditative walk: virtual reality therapy for chronic pain management. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 521-524.

[11] Sidhant Gupta, Dan Morris, Shwetak N. Patel, and Desney Tan. 2013. AirWave: non-contact haptic feedback using air vortex rings. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*. ACM, 419-428.

[12] Felix Hülsmann, Julia Fröhlich, Nikita Mattar, and Ipke Wachsmuth. 2014. Wind and warmth in virtual reality: implementation and evaluation. In *Proceedings of the 2014 Virtual Reality International Conference*. ACM, 24.

[13] Hiroshi Ishii, Sandia Ren, and Phil Frei. 2001. Pinwheels: visualizing information flow in an architectural space. In *Ext. Abst. CHI 2001*: 111-112.

[14] Weina Jin, Amber Choo, Diane Gromala, Chris Shaw, and Pamela Squire. 2016. A virtual reality game for chronic pain management: a randomized, controlled clinical study. In *MMVR*. 154-160.

[15] Lynette A. Jones, and Hong Z. Tan. 2013. Application of psychophysical techniques to haptic research. *IEEE transactions on haptics* 6. IEEE, no. 3 (2013), 268-284.

[16] Terry K. Koo, and Mae Y. Li. 2016. A guideline of selecting and reporting intraclass correlation coefficients for reliability research. *Journal of Chiropractic Medicine* 15, no 2, PMC, 155-163.

[17] Yuichiro Kojima, Yuki Hashimoto, and Hiroyuki Kajimoto. 2009. A novel wearable device to present localized sensation of wind. In *Proceedings of the International Conference on Advances in Computer Enterntainment Technology*. ACM, 61-65.

[18] Sandip D. Kulkarni, Charles Fisher, Eric Pardyjak, Mark Minor, and John Hollerbach. 2009. Wind display device for locomotion interface in a virtual environment. In *EuroHaptics conference, 2009 and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. World Haptics 2009. Third Joint*. IEEE, 184-189.

[19] Sandip D. Kulkarni, Mark A. Minor, Mark W. Deaver, and Eric R. Pardyjak. 2007. Output feedback control of wind display in a virtual environment. In *2007 IEEE International Conference on Robotics and Automation*. IEEE, 832-839.

[20] Jaeyeon Lee, and Geehyuk Lee. 2016. Designing a non-contact wearable tactile display using airflows. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. ACM, 183-194.

[21] Anke Lehmann, Christian Geiger, Bjorn Woldecke, and Jorg Stocklein. 2009. Poster: Design and evaluation of 3D content with wind output. In *IEEE Symposium on 3D User Interfaces (3DUI 2009)*. IEEE, 151-152.

[22] Benjamin Long, Sue Ann Seah, Tom Carter, and Sriram Subramanian. 2014. Rendering volumetric haptic shapes in mid-air using ultrasound. *ACM Transactions on Graphics (TOG)* 33, no. 6 (2014), 181.

[23] Weiquan Lu, Henry Been-Lirn Duh, Steven Feiner, and Qi Zhao. 2014. Attributes of subtle cues for facilitating visual search in augmented reality. *IEEE Trans. Vis. Comput. Graph. 20*, 3, 404-412.

[24] Mitsuru Minakuchi, and Satoshi Nakamura. 2007. Collaborative ambient systems by blow displays. In *Proceedings of the 1st international conference on Tangible and embedded interaction*. ACM, 105-108.

[25] Taeyong Moon, and Gerard J. Kim. 2004. Design and evaluation of a wind display for virtual reality. In *Proceedings of the ACM symposium on Virtual reality software and technology*. ACM, 122-128.

[26] Jason D. Moss, Jon Austin, James Salley, Julie Coats, Krysten Williams, and Eric R. Muth. 2011. The effects of display delay on simulator sickness. *Displays* 32, no. 4 (2011): 159-168.

[27] Omar Mowafi, Mohamed Khamis, and Wael Abouelsaadat. 2015. AirDisplay: Experimenting with air flow as a communication medium. In *Human-Computer Interaction*. Springer, 316-323.

[28] Multicomp. http://www.farnell.com/datasheets/1772996.pdf.

[29] Takuya Nakano, Shota Saji, and Yasuyuki Yanagida. 2012. Indicating wind direction using a fan-based wind display. In *International Conference on Human Haptic Sensing and Touch Enabled Computer Applications*. Springer, 97-102.

[30] Mark Nazemi, Diane Gromala, and Mehdi Karamnejad. 2014. Virtual reality as analgesia: an alternative approach for managing chronic pain. *International Journal of Creative Interfaces and Computer Graphics (IJCICG)* 5, no. 2 (2014): 75-86.

[31] Processing. https://processing.org/.

[32] Nimesha Ranasinghe, Pravar Jain, Shienny Karwita, David Tolley, and Ellen Yi-Luen Do. 2017. Ambiotherm: enhancing sense of presence in virtual reality by simulating real-world environmental conditions. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM 1731-1742.

[33] Rajinder Sodhi, Ivan Poupyrev, Matthew Glisson, and Ali Israr. 2013. AIREAL: interactive tactile experiences in free air. *ACM Transactions on Graphics (TOG)* 32, no. 4 (2013): 134.

[34] Daniel Spelmezan, Rafael Morales González, and Sriram Subramanian. 2016. SkinHaptics: Ultrasound focused in the hand creates tactile sensations. In *2016 IEEE Haptics Symposium (HAPTICS)*. IEEE, 98-105.

[35] Yuriko Suzuki, and Minoru Kobayashi. 2005. Air jet driven force feedback in virtual reality. *IEEE computer graphics and applications* 25, no. 1 (2005): 44-47.

[36] Jouke C. Verlinden, Fabian A. Mulder, Joris S. Vergeest, Anna de Jonge, Darina Krutiy, Zsuzsa Nagy, Bob J. Logeman, and Paul Schouten. 2013. Enhancement of presence in a virtual sailing environment through localized wind simulation. *Procedia Engineering* 60 (2013): 435-441.

[37] Craig Wisneski, Hiroshi Ishii, Andrew Dahley, Matthew G. Gorbet, Scott Brave, Brygg Ullmer, and Paul Yarin. 1998. *CoBuild 1998*: 22-32.

[38] Bob G Witmer and Michael J Singer. 1998. Measuring presence in virtual environments: A presence questionnaire. *Presence: Teleoperators and virtual environments* 7, 3 (1998), 225–240.