

Communiplay: A Field Study of a Public Display Media Space

Jörg Müller^{1,2}

Dieter Eberle²

Konrad Tollmar³

¹Alexander von Humboldt Institute for Internet and Society, Berlin

²Quality and Usability Lab, Telekom Innovation Laboratories, TU Berlin

³Mobile Services Lab, KTH - Royal Institute of Technology, Stockholm

ABSTRACT

We present Communiplay, a public display media space. People passing by see their own contour mirrored on a public display and can start to play with virtual objects. At the same time, they see others playing at remote displays within the same virtual space. We are interested whether people would use such a public display media space, and if so, how and why. We evaluate Communiplay in a field study in six connected locations and find a remote honey-pot effect, i.e. people interacting at one location attract people at other locations. The conversion rate (percentage of passers-by starting to interact) rose by +136% when people saw others playing at remote locations. We also provide the first quantification of the (local) honey-pot effect (in our case it raised the conversion rate by +604% when people saw others playing at the same location). We conclude that the integration of multiple public displays into a media space is a promising direction for public displays and can make them more attractive and valuable.

Author Keywords

Public Displays; Media Space; In-the-Wild Study

ACM Classification Keywords

H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

INTRODUCTION

This paper brings together the research areas of public displays, playful interaction and media spaces. Media spaces connect people in remote locations in a very casual way and are generally assumed to be effective “social catalysts” [11] by encouraging people to initiate and sustain interaction when they otherwise would not. The majority of media space research has addressed work oriented communication (e.g., [16, 19]). Some artistic media spaces have been successfully used, but less researched, in public settings (e.g., [6, 8]).

We know that public displays can be very effective social catalysts by creating a “sociable buzz” around them locally, this is known as the “honey-pot effect” [3]. For this paper we

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org. CHI 2014, April 26 - May 01 2014, Toronto, ON, Canada Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-2473-1/14/04\$15.00.<http://dx.doi.org/10.1145/2556288.2557001>



Figure 1. In the Communiplay system, screens were connected in a public display media space. People could play with virtual objects, and people playing at one location could play with people at other locations. When others were playing at a remote location, this often attracted people to look at and play with the screens (the *remote honey-pot effect* raised the conversion rate by +136%).

more specifically define “honey-pot effect” as (1) the increase of the conversion rate (percentage of passers-by starting to interact) [13] and (2) the increase of the interaction duration, when others are already interacting with a display. While in particular (1) has been observed qualitatively in a number of studies (e.g., [3, 14]), we are not aware of any quantification of this effect. Further, we were interested whether this effect also works remotely in public display media spaces, i.e., whether there also is a “remote honey-pot effect”. To explore these ideas further we designed a public display media space, Communiplay, where people can see their own contour mirrored on a public display and can play with virtual objects. At the same time, they see others playing at remote displays illustrated by color-coded silhouettes (see Figure 1).

We describe a public deployment of Communiplay done in August 2013, where we have studied the following research questions: (1) Is there a remote honey-pot effect? (2) Can we quantitatively validate the existence of a (local) honey-pot effect? (3) Does the honey-pot effect also increase interaction durations? (4) Do more local or remote users increase the honey-pot effect?

Our main findings are:

1. We observed a remote honey-pot effect. The conversion

rate increased by +136% when passers-by saw others playing at remote locations.

2. We could quantitatively validate the (local) honey-pot effect. The conversion rate increased by +604% when passers-by saw others playing at the same location.
3. Other users and passers-by (both locally and remotely) also significantly increased the interaction durations.
4. More users (both locally and remotely) strengthened the honey-pot. The more other users (to a maximum of about 5), the higher the conversion rate and interaction duration.

We were also able to observe a multitude of interactions where people were waving, peeking and poking each other through the silhouettes. We also noted that for the short glances that Communiplay offers to passers-by it was hard to fully understand that the silhouettes on screen represented others playing at remote displays. Moreover, using different backdrops that showed the remote location or adding a location label on the silhouettes had no significant effect. However, being able to see screens in different places led to habits over time to play more regularly and show and talk about the system with friends.

One main implication from this work is that by connecting displays, conversion rate and interaction duration can be increased. Simply put, more connected screens, up to a certain limit, boosted the conversion rate even more. We believe that this generalizes to other public displays, although this needs to be validated in future studies.

We will explain the motivation and implications more concisely in the discussion but first we will talk more about related work, and then describe our study and results in detail to return to this discussion.

RELATED WORK

Playful Public Displays

A number of public display installations have used games as a means to attract people. For example, Flashlight Jigsaw [4] is a public display game played through controllers tracked with a Vicon system. It was designed to enable both multi user and single user play. Most players preferred to play with others than playing alone. Hence we also believe that when displays at different locations would be connected in a media space, this would give people much more opportunity to play together. Other examples are the Red Nose game [15] and Polar Defence [5].

Chained displays [18] presented a study of a space-invaders game played with full body gestures. It was shown that imitation behavior is very strong in teaching people how to play the game. Because the contour of players was shown on the screen, players also imitated what they saw the contour of others do (when both had shared focus). Connecting displays at different locations in a media space, as proposed in this paper, multiplies the opportunities for imitation and thus learning.

Within the art community public displays have been used for a long time to invoke interaction and reflection in playful

ways. Starting with Myron Kruger's seminal work on VideoWall [12], many other works have also used silhouettes in a similar way as Communiplay.

Increasing Interactions with Public Displays

Public displays are electronic displays installed in public spaces. Due to falling display prices, they are rapidly becoming ubiquitous in urban areas and soon many public surfaces may become digital and interactive. They have shown to be able to make people have a lot of fun together [18, 14, 4]. Because of their public nature, they are also available for anybody regardless of their access to technology otherwise (thus, bridging the digital divide). Finally, they can be useful for sparking discussions, reflection, and participation [17, 7].

However, in practically all deployments of interactive public displays, achieving a sufficient number of interactions with the displays was a large concern. Huang [10] investigated how many people passing by actually looked at digital displays installed in different cities. She found that the number of glances was very low, and moreover that glances were often very short.

Brignull [3] deployed Opinionizer at a party, a display where users could type in their opinions on a given topic. He observed the honey-pot effect, i.e. when somebody is already interacting with the display, this draws considerable attention and others start to interact, too. The honey-pot effect has since been noted in a large number of studies [18, 14, 20]. In all these deployments, it was observed that people already interacting with a display attract passers-by to stop, look, and also start interacting. In this paper, we investigate whether the honey-pot can be leveraged to span different locations, i.e. do people interacting at one location or time also attract people at other locations when shown on the screen?

Media-Space

Media spaces are electronic settings in which groups of people can work together, even if they do not reside in the same place or are not present at the same time. In a media space, people can create real-time visual and acoustic environments that span physically separate areas. They can also control the recording, accessing and replaying of images and sound from those environments [16]. Media spaces are particularly good for informal communication between remote places [19].

A major focus of media spaces has been on work settings, although some examples of others, in particular in public art, exist. Hole-in-Space [8] and Bill Fontana's work on relocation of urban soundscapes [6] are among the most prominent. However, most of these art installations have not been researched from an HCI / CSCW perspective.

Other examples are Sports over a distance [1] that extended media spaces with physical exertion, and Palimpsest [1] that provided a large shared projection installed at an art festival. Telemurals [11] is also a more recent system that manipulated the video to anonymize users and overlay them within the same virtual space similar to Communiplay. The system was intended as a social catalyst to initiate and sustain interactions within and between two remote places.

Connected Public Displays

We are aware of only two projects connecting multiple public displays into media spaces. In [9], two locations were connected with a video tunnel. Users were shown with their silhouette overlaid over frozen ice graphics. Via waving, users could melt the ice and see a color image of the other location. It was shown that a symmetric ice melting, where both sides could always see the same area, was preferred by users.

In the screens in the wild project [2], four touchscreens were connected to each other. On the bottom of the screen live video feeds from all four locations were shown. One of the applications on the screen was SoundShape, which allowed collaborative live music making. The authors observed interaction of users with people at the remote location, such as imitating movements, dancing together, or virtual sharing of food by showing it to the camera.

COMMUNIPLAY

Shared Interaction

Communiplay consists of a number of connected large screens. People see their own silhouette mirrored on a public display in front of them and can play with virtual objects (cubes or balls). These objects can be kicked and punched with the silhouette and behave according to simulated physics. Opposed to normal games there are no scores to collect, so the main driver is simple to bounce and juggle around with the objects.

At the same time, the players also see others playing at remote displays within the same virtual space illustrated as silhouettes (Figure 2), with the same virtual objects so the physics are shared among all screens. Each location has its own color for the silhouettes of the user to distinguish their location. Additionally a text label telling the location is shown on each silhouette. The size of the silhouette shown on the display is determined by the distance of the users to the displays, i.e. like a mirror.

On the backdrop of the screen the users will get some additional clues about the connected places in an animated banner that is rendered on top of an abstract whiteish image. We also tested to use a "live" caption as background to emphasize how the Communiplay system was connected (Figure 2, e and f) and we will soon discuss how these configurations and arrangements of the systems affected the use of the system. However, the system did not support any additional back-channeling modes such as audio feedback or speech between the remote locations since we wanted to first study how the game and silhouettes facilitate the shared interaction.

System

At each display a depth camera is mounted to scan the environment for passers-by. The software runs under Ubuntu linux on notebooks that are attached behind the screens. The screens are connected to a server that runs the game engine and keeps track of the shared interaction of the virtual objects. The logging is done on the server-side as well.

The implementation for both client and server was done with Java with the help of some external libraries like JavaCV, a

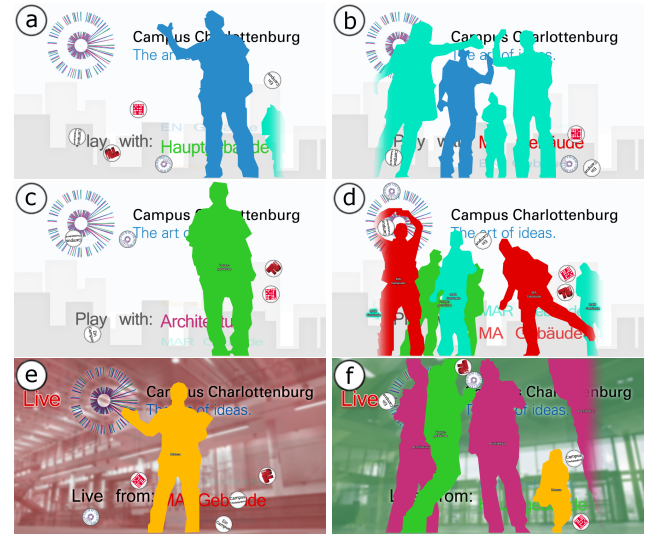


Figure 2. Location labels: (a,b) Users were shown their silhouette without labels. (c,d) A location label was shown on each user. (e,f) A "live" caption was added and the background rotated through images of each of the locations.

wrapper for the OpenCV implementation that focuses on real-time image processing which is needed to handle the camera input. For the access of the depth camera OpenNI and NITE were used.

All silhouettes are based on the user masks provided by OpenNI. These data are simplified with OpenCV to reduce the fine user masks to polygons. This reduces the network load to a few sets of points and Netty provided support for the networking part.

The physics simulation of the balls is realized with a simple contour collision algorithm. With the use of FFmpeg anonymous depth camera pictures are saved for analysis of logs. The whole scene is rendered with OpenGL with the LWJGL library. The physics calculation is done both on the server-side and the client side so all users interact with the same objects, but interaction continues in case of network outages.

PRESTUDY

The objective of this study was to explore audience behavior with a public display media space and to compare different grades of anonymity (contour vs. mirror image, see Figure 3). Three displays were deployed for two days on three different floors of a research laboratory. One display was installed in the cafeteria, one in the lobby next to the reception, and the third in a hallway close to workplaces. During the first day,

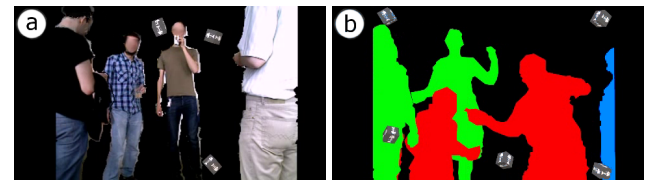


Figure 3. Prestudy conditions. (a) Mirror condition, (b) contour condition.

the mirror (colored rgb) image of the users was shown (image condition). On the second day, only the contour of users was shown (contour condition). We observed usage while sitting in an inconspicuous location close to the screens. We also collected anonymous depth videos of usage.

Findings

Remote Honey-pot Effect

Although we did not quantify this, when somebody played at one of the locations, there seemed to be a higher probability that people passing by the other screens would stop and start to play. Consequently, there were often people in different locations playing simultaneously. Sometimes, this went so far as to create a kind of overcrowding. There were so many people in front of the screen that they actually could not interact with the virtual objects anymore and that not all of them were visible. It was however unclear if this overcrowding would prevent passers-by from starting to interact or make people leave the screen.

Playing the game vs. playing with the mirror images

When people interacted with the screens alone, naturally they played mostly with the virtual objects. One behavior we saw repeatedly was people trying to grab and pick up the objects. This, however, was usually unsuccessful, because the resolution of the camera was too low, and grabbing as a gesture was not supported by the application. Another common behavior was boxing objects and sometimes fighting for control over the objects with people at the other location. When there were people at other locations however, users interacted much more with the other users than with the objects. In the contour condition, people played with each others contour, e.g., by running back and forth and trying to hide each others contour and creating funny poses. Especially in the image condition, but also in the contour condition, users mostly waved towards each other and looked at each other. Some people just played with their contour, for example by dancing, even when they were alone.

Ghost effect

One funny effect occurred when people played with the screen and somebody else entered the screen at a different location. People would then turn around, apparently because they expected the other person to be behind them. When they saw that there was nobody, they would often be surprised, look at the screen again, and then apparently realize that the other person probably was at a different location.

Discussion

The most interesting aspect for us was the remote honey-pot effect. If people interacting at one location could attract others to interact at a different location, the connection of different displays to each other could make public displays much more attractive. The flip side of the coin was the overcrowding. Especially when many screens are connected to each other, there would be a high probability that too many people would be on the screen to enable everybody to interact. It is unclear however whether such an overcrowding would be negative or positive. With the real honey-pot effect, usually the bigger the crowd the more people it attracts.

Although the mirror image has very interesting properties (users can recognize each other) we decided to continue with the contour condition. It is much easier to deploy for many public locations. The provided anonymity makes privacy issues less of a concern. In some countries (like Germany) it is also simply not legal to distribute image recordings of passers-by in public spaces without their consent. Finally, it is also more playful and technically simpler.

FIELD STUDY

The main objective of the field study was to investigate usage of a larger public display media space in a more public setting. In particular, we wanted to confirm and quantify the honey-pot effect, and investigate the existence of a remote honey-pot effect. That is, we wanted to confirm that more people passing by would interact with the displays when 1) others interact with the display at the same location, or 2) others interact at a different location. We also wanted to confirm that the interaction duration would increase in these two conditions. Additionally, we wanted to compare the effect of real people who are streamed live from a different location to recordings of users who have interacted before. For this reason, recordings of users who interacted for more than 10s were taken and played back at random intervals when nobody else was on the screen. We conducted a two-phase deployment. In a first phase of three days, we collected observations of system and user behavior and iteratively improved the system. In the second phase of two weeks, we collected data without changing the system.

Deployment

We deployed Communiplay in six different buildings of the Technical University of Berlin, Germany for two weeks in August 2013. The locations are close to the entrance of the main cafeteria (CAF), the main building (MAIN), in the architecture building (ARCH), in the electrical engineering building (EE), in the mathematics building (MATH) and the computer science building (CS). The chosen spots here were by the main traffic path to maximize the number of passers-by and situated so that passers-by would see them and naturally walk through the interaction area to maximize inadvertent interactions [14]. The typical passers-by are naturally students but off semesters like August the percentage of faculty are rather high. Moreover the university campus is located in the middle of Berlin and hence also visited by many others, like tourists and people just walking across the campus.

Optimization Phase

On the technical side, we installed much stronger external WiFi antennas, considerably reducing networking problems. In the pre-study we had the cameras installed immediately below the screens to improve the perspective. This posed two kinds of problems that occurred in the fieldstudy because of the different space. First, for far away users they would hover “in the air”, not being able to kick the objects. Second, patches of the floor were often recognized as false users when a user had passed by. A relocation of the cameras to the floor below the screens solved both problems. We also collected much more data than expected, causing problems

with the RAM of the server. Resorting to the hard-drive for logging data helped. We also changed the physics simulation to continue simulating locally when the connection to the server was lost for more than a few milliseconds. Also, the rendering of the local silhouette was made independent of the server connection, strongly reducing latency. We noticed log data drift issues for the different logs kept locally and implemented logging on the server instead. We changed one screen to a different location because of networking issues and more people passing by in the new location. From interviews we quickly learned that people had difficulties understanding that the screens were connected. To address this, we added a hint on the screen stating which locations the screen was connected to. To lower the impact that motion alone of remote users would have compared to no users, we made the location hint move so we get one constantly moving item on the screen. We also noticed that it was considerable effort checking the state of the screens, and getting a feeling what was going on. Therefore, we installed a viewer close to our workspaces where we could permanently see what was going on on the screens, and also whether any screens were offline. We learned that it was distracting for us not to know whether users shown on the screen were real users or recorded users. We added small color pixels on the side of the screen to indicate recorded users for us.

Design

In order to investigate the impact of other users on user behavior, we used a quasi-experimental design in the sense that we did not control the independent variables (conditions). Instead, users were assigned to conditions based on the behavior of other users and the system. We distinguish six different conditions, when a user is passing by a screen:

1. Other users play in front of the same screen (local honey-pot)
2. Other people pass by the same screen
3. Other users play in front of screens at other locations (but not at the same, remote honey-pot)
4. Other user pass by another screen
5. Recordings of users who played at different times are shown (but nobody currently plays, recorded honey-pot)
6. Nobody is shown on the screen (baseline).

Condition five (5) was generated when nobody else was on the screen. More specifically, at a random time between 30 and 90 seconds after the last passer-by, a recorded session from a past interaction was started. Throughout this paper and in the video, recorded users are highlighted (black). However, in the experiment, they were indistinguishable from real users.

In order to assign each user to one of these quasi-conditions, a time frame of 10s before the user entered the tracking area of the camera was considered. This was done because from our experience of the honey-pot effect, people are often attracted by others who played before them, but only approach the screen after the others have left.

Dependent variables were percentage of passers-by who interact (conversion rate) and duration of interactions. Our hypotheses were (1) existence of a local honey-pot effect: the conversion rate increases when people already interact at the same screen, as does the interaction duration, and (2) existence of a remote honey-pot effect: The conversion rate increases when people who interact simultaneously at other locations are shown on the screen, as does the interaction duration.

We were also interested to test how the representations of the silhouettes and background affected the interaction. The way we tested this was that we also added two additional variables. Firstly we added location labels to the silhouettes, which were toggled on and off in a interval of 15 minutes. For the second half of the deployment, we compared the location labels with a different background from an abstract design to a "live" caption background from the different locations in a interval of 15 minutes (Figure 2).

Data Analysis

We logged the anonymous depth video from all displays during the entire duration. In addition, we logged user behavior as determined by OpenNI (entering and leaving users). Further, we logged screen capture during the entire deployment. We followed a semi-automatic video coding approach. We used the OpenNI logging to automatically determine all situations where somebody was in front of the display. We then manually decided for each person whether the detected person was passing-by, spectating, interacting, or noise.

If the person was interacting, we noted the beginning and the end of the interaction. We further categorized the interaction within the categories of playing with objects, waving / gesturing towards others, playing with contours (hiding others etc.), making funny shapes with contours, imitating others at same location, imitating others at other location, and other.

The data was annotated by four raters. A portion of the videos were annotated by all raters to determine inter-rater reliability. Inter-rater reliability was substantial (Fleiss' Kappa = 0.75 for interactions, and 0.66 for spectating). We also manually marked situations where something interesting or unusual happened. Our annotations included a pointer to the situation in the video, so we could easily go back to the original data to verify our annotations.

We also conducted semi-structured interviews with users and spectators. In particular, we asked them to describe the displays, asked whether they had communicated with people at the other location and whether the others had responded, what they thought where the people at the other location were and whether they felt connected to them. The interview data was anonymized and put into a shared form that was analyzed by two researchers.

Results

Over the two week field test we have been able to record and analyze 30888 people by-passing the screens and 1234 interactions with the screens. The average conversion rate over all conditions was 4%. According to our experience, this is

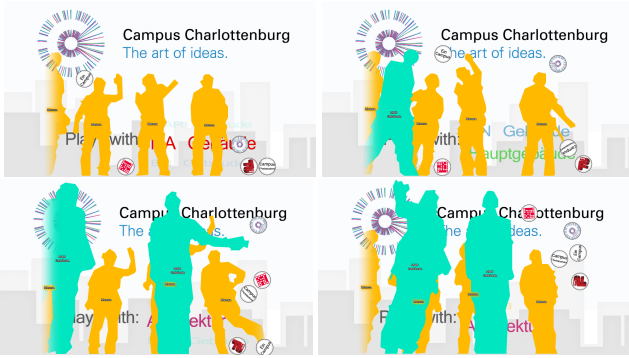


Figure 4. Remote honey-pot effect: Users attracted other users from another location.

in line with other public display deployments and industry values. We will now breakdown the data and show the key results.

Remote and Real Honey-pot Effect

A Kruskal-Wallis test revealed a significant effect of condition on conversion rate ($\chi^2(5) = 788.5, p < .01$). A post-hoc test using Wilcoxon tests with Bonferroni correction showed significant differences between the Baseline and all conditions except the Recorded User, between Remote Passer-by and Remote User, and between all conditions and the Local User ($p < .05$), see Figure 5.

So, we found a remote honey-pot effect. Observing people interacting at a different location (condition Remote User) increased the conversion rate by 136% compared to the baseline (see Figure 4 for an example). Recorded users increased the conversion rate only by 38% (differences to Baseline and Remote User were not significant). We also found, what we believe to be the first experimental proof, of the honey-pot effect for public displays. When others were already interacting at a screen (condition Local User) the conversion rate increased by +604% (see Figure 6 for an example).

Honey-pot for Interaction Duration

We found that if there were remote passers-by, interaction duration was significantly longer than if not (mean 33.3s vs.

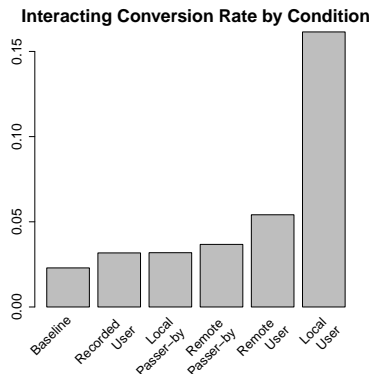


Figure 5. Conversion rate by condition. The remote honey-pot (Remote User) increased the conversion rate by +136% and the local honey-pot (Local User) by +604% compared to the baseline.



Figure 6. Typical example of the honey-pot effect. Two users interacted with the screen, which attracted the attention of a passer-by (1). As the two previous users left (2) the passer-by approaches the screen (3) and started to interact (4).

9.3s, t-test, $t(1176) = 11.4, p < .01$). If the remote users interacted, interaction duration was even longer (mean 53.7s vs. 33.3s, t-test, $t(386.5) = 3.8, p < .05$). The same was true locally. If there were local passers-by, interaction duration was significantly longer than if not (mean 15.7s vs. 30.4s, t-test, $t(817) = 5.5, p < .01$). If the other local users interacted, interaction duration was even longer (39s vs. 30.4s, t-test, $t(1098) = 2.72, p < .01$). We also found that if local users or passers-by left, it was significantly more likely during the next 5s that another user left than otherwise ($\chi^2(1) = 435.2, p < .01$ for users, and $\chi^2(1) = 53.9, p < .01$ for passers-by), while this was not the case for remote or recorded users leaving.

More Users Strengthen the Honey-pot

Finally, we found a significant effect of the number of remote users on the remote honey-pot effect (see Figure 7). The more users were interacting remotely, the higher the conversion rate (up to about five remote users). A Kruskal Wallis test revealed a significant effect of number of remote users on whether passers-by interacted ($\chi^2(10) = 916, p < .01$). A post-hoc test using Wilcoxon tests with Bonferroni correction showed the significant differences between 0 and all except 10 users, 1 and all except 8 and 10, 2-5 and 3-5 remote users ($p < .05$).

This was also true for the local honey-pot effect. The more users interacted, the higher the conversion rate. A Kruskal Wallis test revealed a significant effect of number of local users on whether passers-by interacted ($\chi^2(7) = 806.7, p < .01$). A post-hoc test using Wilcoxon tests with Bonferroni correction showed the significant differences between all

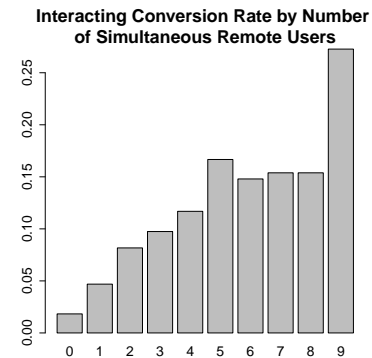


Figure 7. The remote honey-pot effect grew stronger with more remote users.

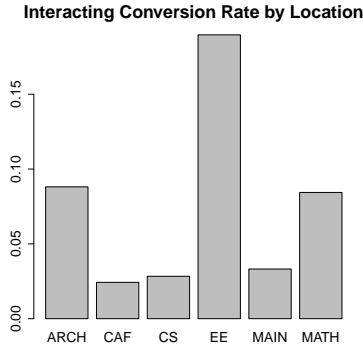


Figure 8. Conversion rate by location.

pairwise comparisons except 3-4, 4-7, 5-7 and 6-7 ($p < .05$).

Last, with one-way ANOVA, we found a significant effect of the number of other locally interacting users on interaction duration ($F(7,1226) = 11.7, p < .01$), and of other remotely interacting users on interaction duration ($F(9,1224) = 17.2, p < .01$). The interaction duration increased with the number of other users.

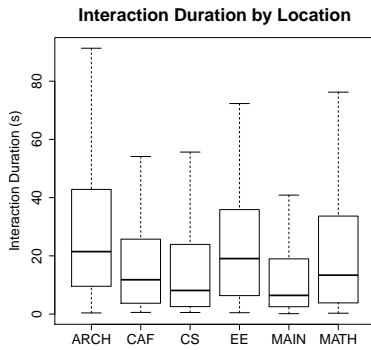


Figure 9. Interaction duration by location.

Differences between Locations

As described above, in one of the conditions we put nametags on the silhouettes. We thought that this would make it clearer with whom you were interacting but this had no significant impact on conversion rate or interaction duration. However we found significant differences in usage between locations, both regarding conversion rate (Figure 8) and interaction duration (Figure 9). A Kruskal-Wallis test revealed a significant effect of location on conversion rate ($\chi^2(5) = 697, p < .01$). A post-hoc test using Wilcoxon tests with Bonferroni correction showed the significant differences between all locations comparisons except MAIN-CS, ARCH-MATH, and CS-CAF ($p < .05$). With one-way ANOVA, we found a significant effect of location on interaction duration ($F(5,1228) = 4.4, p < .01$). A Tukey's pairwise comparison revealed significant differences between locations ARCH-MAIN, CS-EE, and ARCH-CS ($p < .05$). Hence a strong pattern that might explain the rather large variances of duration of the interaction would be the different character of space for each screen.

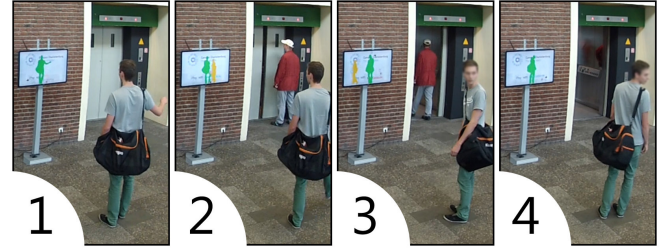


Figure 10. Many users had difficulty understanding that the screens were connected in a media space, despite our efforts of communicating that on the screen. When they saw somebody pass at a remote location, sometimes they would turn around, apparently in the expectation that the other person would be behind them (*ghost effect*).

Our study design did not allow us to collect quantitative data for further analysis but we will talk more about the different types of places in the discussion.

Interaction with Communiplay

Video Observation

From the analysis of the recorded depth video we have been able to observe a variety of known and also novel interaction patterns with Communiplay. People played mostly with the objects, but also communicated with remote users, e.g., by waving, making funny shapes, playing with the silhouettes, and mimicing local and also remote users. The frequency of interactions is given in Figure 11.

One recurring effect that we observed rather frequently was that when somebody passed by a remote screen, people turned around, expecting that the person was behind them. Because of the frequency of this observation, we called this the “ghost effect” (examples are shown in Figure 12 and Figure 10, and the effect was also observed in the prestudy). Although people from different locations were drawn in different colors, and had location labels drawn over their bodies, was apparently was not enough of a hint that they were at a different location. Sometimes, the shared interaction between the different locations was also experienced as annoying. In particular, when people tried to perform more complex interactions, like picking up an object, remote by-passers could

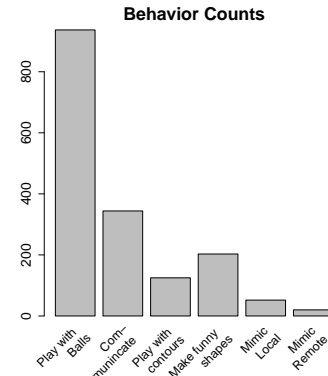


Figure 11. Number of occurrences of different behaviors. Users spend most of the time playing with balls, followed by communicating (e.g., waving), making funny shapes, and playing with the contours.

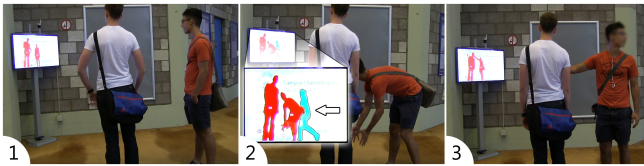


Figure 12. One downside of the connection of the screens occurred sometimes when people tried to interact with the display, like this person trying to pick up a ball. Then, when a remote person walked through the screen, they would often kick the balls away (also an example of the ghost effect).

disturb the interaction by causing all objects to fly around (see Figure 12).

When users tried to communicate with each other, this was mostly through waving (Figure 13). Users both waved to say hello as well as goodbye. We also observed a number of cases of people imitating others who interacted in the same location or remotely.

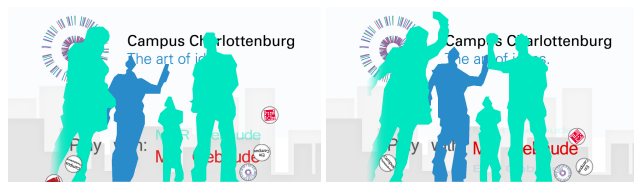


Figure 13. People waving each other when entering the display.

One example of playing with silhouettes is given in Figure 14. Sometimes, either because a touch was misinterpreted or because one user started punching each other, a (remote) fight broke off. People performed surprisingly acrobatic fight moves, like high kicks. Most of the playing with silhouettes was however more peaceful, like people poking each other or shaking hands.

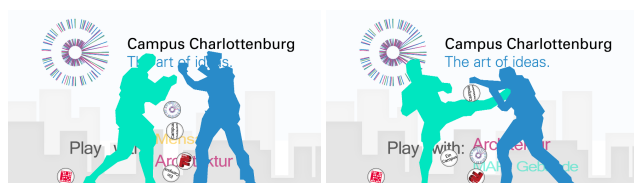


Figure 14. Two users in different locations fighting each other.

We also observed a number of interesting interactions with recorded users. During the deployment, the recorded users were drawn like remote users, so there was no easy way for users to distinguish between them (but in the figures, we highlighted recorded users in black). For example, Figure 15 shows a user mimicking two recorded users. Figure 16 shows a user trying to interact with a recorded user through poking, punching, and playing with balls.

Interviews and Live Observation

We also did live observations and semi-structured interviews with about 100 persons, and in 55 cases they also filled in a small survey. We would like to stress that it is very hard to stop people in public settings that are just passing by, so most of our respondents either stopped for looking at or interacting with the system.



Figure 15. A user mimics a recorded user (recorded user highlighted in black).

We were a bit surprised that apparently it was very difficult for users to understand that the screens were connected. Probably because most existing media space research was done in relatively tight groups over prolonged time, we were not aware of reports of users having difficulty to understand that the screens are connected. 32% of the respondents noticed in forehand that it was other real persons on the screen and that the silhouettes where not just a screen decoration, but the remaining 68% did not explicitly notice this and were primarily stopping for interacting with the objects. Possibly because of the physical design of the screen stands and the placement in entrances, many users expected that the screens would be information screens. Some even expected them to be touch screens (which they were not). One interviewee said that “the balls on the ground looks like the buttons on a touch screen, you should explain the game better”. When people at remote locations were passing by, they were often mistaken for video recordings or people passing behind the user locally. Some people said that “[the] connectivity sense of the other place [could be improved]”, or “there should be better evidence that you can play with others”, while we thought that we had already included quite explicit hints, which were apparently not noticed by many users.

Over time, more users learned that the screens were connected. Many said that they “got it” when they saw the screens in other locations, or when they saw somebody passing by remotely. Some people also reported that they developed a habit to wave every time they passed by the displays (mostly without slowing down). They stated that they hoped somebody would wave back, and sometimes somebody would. Some cafeteria employees reported that they went to the screens for each of their breaks to play a bit, and sometimes someone at another location would play with them.

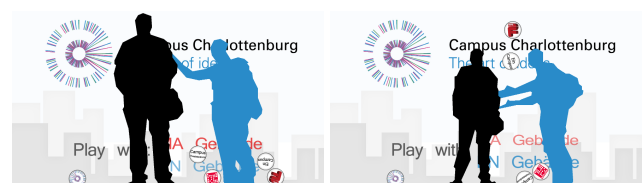


Figure 16. A user touches and tries to interact with a recorded user (recorded user highlighted in black).

Another common pattern was that one person of a group stopped and demonstrated the system to some friends or colleagues. During the interview, the demonstrating person often took a leading role and explained that they had used it before or have seen it at some other location, and now wanted to show it to their friends.

When we asked why they stopped the most common reply was that the screens were attractive visually, like: “[I] just stopped for a short glance, saw something interesting”, or “[I stopped] to look and play”. Furthermore the game was experienced rather fun to play with but comments also indicated that the overall experience was rated higher than the actual game.

Many comments focused around ideas how to improve the game and how to make the communication more direct, i.e. “perhaps a score list of the sites”. Also people in groups were more likely to make a quick stop than single persons especially in busy areas. Overall we noted some unease by single persons to interact in front of the display. They preferred to first spend a moment or two observing what was going on in the local proximity.

When we asked what the silhouettes represent, most find it hard to put in words what they experienced and the replies become rather vague. Many also avoided to speculate and said that they had not fully understood the interaction. However when asked what they focused on during the interaction: 36% told us they played with the objects, 32% observed other silhouettes, 25% played with their own contour. Only 7% told us that they played with contours of remote people.

DISCUSSION

By primarily studying the first seconds of interaction (getting people to stop and start interacting) we have learned more about the use, what attracts the usage of Communiplay and the key factors that impact these findings. We will now discuss some of the core results in more detail.

Honey-Pot

The key finding in this study is that we could show the existence of a honey-pot effect. The strongest honey-pot effect was local presence, but also remote presence on the screens had a significant honey-pot effect.

Possible reasons may vary from quite simple, like the fact that motion attracts attention, over more complex, that people exhibit imitation behavior, or are attracted by crowds, to high level reasons, like that multiplayer games are more fun than single player games, or that playful communication is more interesting than just playing.

It has been shown that moving, especially looming, stimuli attract visual attention [13]. Remote people shown on the screen cause motion. Thus it may be that the conversion rate increases simply because this motion attracts more attention. However, we had partially addressed this by including permanent motion on the screen, and this would not explain the apparent difference between remote and recorded users.

One caveat with our study design is that the conditions are not randomly distributed over time of day. It could therefore be that the conversion rate is simply different on different times of day, e.g. higher at noon, and that some conditions (e.g., other simultaneous users) occur more frequently e.g., at noon, and that therefore it would appear as if the conversion rate would be higher with simultaneous users. As a precaution, we used the conversion rate as the dependent variable, which

is much more stable over time of day than, e.g., number of users per hour. As a second precaution, we conducted the analysis separately for morning, noon, afternoon, and night, and the effects seem to be stable over these different periods.

From the interviews and observations combined with our data we can now speculate around the difference between local and remote honey-pot effect. One reason could be that people do not fully understand the remote representation and hence the effect becomes weaker. Another reason might be that handling personal space is easier done in real physical space and people feel too close to each other on screen, or do not want to disturb/interrupt others.

Interestingly, recorded users seemed to be less effective than remote users (although not significant). However, our experiment did not allow us to determine the reason for these effects so we can only speculate that the lack of shared interaction had more impact than was unrevealed in our observations and interviews. Moreover, having two people at two different locations at overlapping times is less frequent so even if we have not been able to fully demonstrate the benefit of using recorded scenes we still think this an interesting strand for further explorations.

As stated earlier, the duration of the interaction was not significantly different across the different conditions. However differences in duration of the interaction were significant between the different places and we will discuss this next.

Type of place

We observed a clear difference in conversion rate but also interaction duration between the six different locations that were used in the fieldstudy. From conversion rates and interaction durations, it would look as three (CAF, CS and MAIN) would be the least successful locations.

However, looking at the total number of interactions, it can be seen that these locations actually have the highest number of interactions, but then lower conversion rates and shorter interaction durations.

In future work it would be interesting to look more into how the flow in a place affects the conversion rate and how to optimize the different combinations of honey-pot effects we discussed here, i.e. local and remote, due to the nature of the place.

Playful remote interaction

People did not only interact with the shared objects in Communiplay but also used a whole range of behaviors to more playfully interact and communicate with persons on the screen, i.e. by waving, peeking and poking, and mimicking behavior.

The “ghost effect” was an example of people not understanding that other users were at a different location. It would be interesting to further explore how different levels of temporariness and abstractions affect the conversion rate and duration of the interaction, as well as the experienced interaction and the awareness of the remote persons, and how this comes together. Moreover, different forms of representations

could also be further explored. As noted earlier people just passing by sometimes ruined the experience for people interacting with Communiplay and it might be a good idea to differentiate between passers-by and people interacting with Communiplay, and e.g., not letting people just passing by to interact with the objects.

CONCLUSION

In this study we have learned that the integration of multiple public displays into a media space is a promising direction for public displays and can make them more attractive and valuable. We observed a local and a remote honey-pot effect. Further, both of these effects also increased the interaction duration. Finally, more users strengthened the honey-pot effect (both locally and remotely, and both for conversion rates and interaction durations). The over-crowding effect that showed in the pre-study was not present in the public deployment. We believe that these effects generalize to other kinds of public displays, although this needs to be validated in future studies.

We learned that the playful interaction with remote people offered by our system works even if the communication is very temporal and abstract. Moreover, people did not only play with the shared objects but also used many other means to more playfully interact and communicate, i.e. by waving, peeking and poking, and mimicing behavior.

The main implication from this project is that public displays and media spaces can be integrated effectively. The local honey-pot created by public displays translates to a remote honey-pot when displays are connected into a media space. Thus, such a public display media space can become an effective social catalyst, both locally and remotely.

ACKNOWLEDGEMENTS

We particularly want to thank Constantin Schmidt for tremendous support during the deployment, and Tiare Feuchtnner and Andreas Fender for help during annotation. This work was supported by the ICT Labs of the European Institute of Innovation and Technology.

REFERENCES

1. Agamanolis, S. Designing displays for human connectedness. In *Public and Situated Displays*. Springer, 2003, 309–334.
2. Behrens, M., Fatah gen. Schieck, A., Kostopoulou, E., North, S., Motta, W., Ye, L., and Schnadelbach, H. Exploring the effect of spatial layout on mediated urban interactions. In *Proc. of PerDis '13*, ACM (New York, NY, USA, 2013), 79–84.
3. Brignull, H., and Rogers, Y. Enticing people to interact with large public displays in public spaces. In *Proc. of INTERACT '03*, vol. 3 (2003), 17–24.
4. Cao, X., Massimi, M., and Balakrishnan, R. Flashlight jigsaw: an exploratory study of an ad-hoc multi-player game on public displays. In *Proc. of CSCW '08*, ACM (2008), 77–86.
5. Finke, M., Tang, A., Leung, R., and Blackstock, M. Lessons learned: game design for large public displays. In *Proc. of DIMEA '08*, ACM (2008), 26–33.
6. Fontana, B. The relocation of ambient sound: Urban sound sculpture, 1987.
7. Fritsch, J., and Brynskov, M. 7 between experience, affect, and information: Experimental urban interfaces in the climate change debate. *From Social Butterfly to Engaged Citizen: Urban Informatics, Social Media, Ubiquitous Computing, and Mobile Technology to Support Citizen Engagement* (2011), 115.
8. Galloway, and Rabinowitz. Hole in space, 1980.
9. Häkkinä, J., Koskenranta, O., Posti, M., Ventä-Olkkonen, L., and Colley, A. Clearing the virtual window: connecting two locations with interactive public displays. In *Proc. of PerDis '13*, ACM (2013), 85–90.
10. Huang, E. M., and Mynatt, E. D. Semi-public displays for small, co-located groups. In *Proc. of CHI '03*, ACM (2003), 49–56.
11. Karahalios, K., and Donath, J. Telemurals: linking remote spaces with social catalysts. In *Proc. of CHI '04*, ACM (2004), 615–622.
12. Krueger, M. W., Gionfriddo, T., and Hinrichsen, K. Videoplace - an artificial reality. In *ACM SIGCHI Bulletin*, vol. 16, ACM (1985), 35–40.
13. Müller, J., Alt, F., Michelis, D., and Schmidt, A. Requirements and design space for interactive public displays. In *Proc. of MM '10*, ACM (2010), 1285–1294.
14. Müller, J., Walter, R., Bailly, G., Nischt, M., and Alt, F. Looking glass: a field study on noticing interactivity of a shop window. In *Proc. of CHI '12*, ACM (2012), 297–306.
15. O'Hara, K., Glancy, M., and Robertshaw, S. Understanding collective play in an urban screen game. In *Proc. of CSCW '08*, ACM (2008), 67–76.
16. Olson, M. H., and Bly, S. A. The portland experience: A report on a distributed research group. *International Journal of Man-Machine Studies* 34, 2 (1991), 211–228.
17. Schroeter, R. Engaging new digital locals with interactive urban screens to collaboratively improve the city. In *Proc. of CSCW '12*, ACM (2012), 227–236.
18. Ten Koppel, M., Bailly, G., Müller, J., and Walter, R. Chained displays: configurations of public displays can be used to influence actor-, audience-, and passer-by behavior. In *Proc. of CHI '12*, ACM (2012), 317–326.
19. Tollmar, K., Chincholle, D., Klasson, B., and Stephanson, T. Videocafe - exploring mediaspaces in public places within organizations. *Behaviour & Information Technology* 20, 2 (2001), 101–110.
20. Walter, R., Bailly, G., and Müller, J. Strikeapose: revealing mid-air gestures on public displays. In *Proc. of CHI '13*, CHI '13, ACM (New York, NY, USA, 2013), 841–850.