In-Situ Occlusion Resolution for Hybrid Tabletop Environments

Jan Riemann, Mohammadreza Khalilbeigi, Max Mühlhäuser

In-Situ Occlusion Resolution for Hybrid Tabletop Environments

Jan Riemann, Mohammadreza Khalilbeigi, and Max Mühlhäuser

Technische Universität Darmstadt, Germany
{riemann, khalilbeigi, max}@tk.informatik.tu-darmstadt.de

Abstract. In this paper we explore the use of in-situ occlusion resolution in mixed physical/digital tabletop scenarios. We propose the extension of back-projected tabletops with interactive top-projection to turn the physical object's surface into peripheral displays. These displays are used to resolve occlusion in-situ without the need to use additional tabletop display space and keeping the spatial perception of the occluded objects. We contribute a visualization concept and a set of interaction techniques for in-situ occlusion resolution and easy access to occluded objects. The techniques are implemented in a system named ProjecTop, which is evaluated in an quantitative user study. The study results highlight how top-projection can be beneficially used. We conclude with a set of design implications derived from the study's results.

Keywords. Interactive tabletops · occlusion awareness · hybrid interaction · peripheral displays · multitouch

1 Introduction

Tabletop computers are going to play an important role in future offices due to dropping prices and technological advances. When integrated within the normal office furniture, their table like form factor affords the usage as normal table, e.g. additional materials like printed documents, laptops, books etc. are placed and used concurrently on its interactive surface.

Fig. 1. ProjecTop system visualizing the occluded digital objects

The presence of physical objects on tabletops poses several challenges: The first is that it considerably decreases the screen real-estate for interaction with digital objects. Not only the footprint of physical objects obstructs the display, but also casual arrangement of them on the surface makes the - still visible - display areas hard to use. This becomes even worse for tall physical objects (e.g. high stacks) that shadow even more display area from the user's perspective. A second important challenge is that the physical objects may partially or entirely occlude digital objects, resulting in losing awareness about occluded objects. Further, accessing them becomes cumbersome.

Work is emerging to address parts of these problems: detecting the footprint of physical objects [2] and avoiding occlusion [1]. Other researchers address the problem by providing support for staying aware of and accessing digital objects [6,3,8]. The typical approach is that the occluded objects are transformed into icons or other content-wise reduced forms and visualized next to the occluding object. While this pioneering and motivating work established first design principles to mitigate the problem of physical occlusion in hybrid settings, it is not clear how these approaches extend to more cluttered workspaces where many physical objects consume considerable space on the tabletop display, which is the user's primary interaction surface.

We go beyond existing work by extending the interaction and display spaces to the surface of the physical objects using additional top-projection facilities on a back-projected tabletop system to resolve occlusion. The tabletop provides the primary high-resolution interaction surface for the user to work on while the top-projection provides a secondary class of “peripheral” displays on the physical objects' surface. These peripheral displays are used to solve the aforementioned problems. By using the object's surface as secondary display, we do not need to consume additional space on the tabletop screen. Another advantage is that the spatial layout can be left unchanged, as the visualization can be shown at the occluded object's exact location.

In this paper, we contribute a set of integrated interaction concepts and visualization techniques that leverage the underlying two-class display idea for in-situ occlusion resolution in hybrid environments. The design of our techniques is grounded on an initial user study. All presented techniques are coherently implemented in a fully functional occlusion support system, named ProjecTop, allowing the users to fluidly interact in a realistic hybrid tabletop setting.

Our second contribution is a user study formally evaluating our approach in a real hybrid setting. The goal was to assess the potential of top-projection in hybrid settings in general and for occlusion management in particular. Based on the results, we present design implications for hybrid tabletop systems extended with top-projection. We believe that the results are not only relevant in the context of occlusion but also in general for the use of top-projection to project additional self-contained content on physical objects.

The outline for the rest of the paper is as follows: We start with a review of the related work. We then present an initial user study followed by the design of the ProjecTop system based on the study results. Next, we present the user study evaluating the ProjecTop system and conclude with design implications derived from the study.

2 Related Work

Our work mainly relates to two research fields - namely augmented desktops and physical occlusion - which are discussed in the following:

2.1 Augmented Desktops

Effort has been made to augment paperwork with digital projection: The DigitalDesk [22], FACT [11] and Paper-Top [13] are systems which project onto printed documents, for example allowing the user to interactively add digital content [22,11,13], show animations [13] or do highlighting [11]. The EnhancedDesk [10] uses a projector to display related digital content (e.g. interactive simulations) next to physical documents. However, these systems aim at augmenting paper documents with additional digital facilities and do not take into account self-contained digital objects. Other examples are the Augmented Surface system [16], which allows digital objects to “leave” devices like laptops by means of top-projection on desk and wall surfaces, or Pictionaire [4] which is a projected tabletop system to integrate digital and physical media. Projection on physical objects is used by Pictionaire to facilitate tracing of digital content into sketchbooks by projecting the original digital content on the book. There are also some systems, like PaperLens and spatially aware tangible displays [19,17,18], which use top-projection and passive devices to extend the interaction space from the tabletop's surface to the space above the table. This allows for example browsing three dimensional information spaces in a natural way.

While these systems allow the user to circumvent occlusion by physical objects due to the possibility of flexibly shaped display areas [1], their inherent ability to project on the occluders' surface or by using computer vision techniques to move the display around physical objects, these systems have some drawbacks: They either avoid projecting on physical objects and lose display space where physical objects are placed. Or, they do not take into account physical objects and project unaltered content on them, basically ignoring the presence of objects which possibly renders the information on the object itself unusable. Another common problem to all systems that use top-projection as their primary means of display is that they are prone to occlusion by the users' body parts. The user causes shadowing when body parts like hands or the head come into the projection frustum. Therefore, we use top-projection only as a secondary display for occlusion resolution. The primary content is displayed on the back-projected tabletop display which is not prone to these problems.

2.2 Physical Occlusion on Tabletops

Occlusion has been widely studied as it is a common problem - even in physical-only settings. Iwai et al. developed the limpid desk [5] which allows for ad hoc transparentization of stacks of physical objects to browse them. Underlying objects are projected on the topmost object if the user wants to look at them. However they investigated physical-only settings which have the inherent advantage that the presence of physical objects is more visible due to the height of the objects than it is for digital items. Our work extends this approach to digital items on a hybrid interactive tabletop.
Steimle et al. [20] as well as Khalilbeigi et al. [7] conducted studies to gain an understanding on how physical-only interaction transfers to mixed digital-physical interaction scenarios on interactive surfaces. They found out that concepts common in the physical world, like stacking, are also applied to hybrid settings by the studies' partic-

ipants. Therefore the problem of physical/digital occlusion can be assumed to be an important factor for the usability of such systems.

Freeman et al. [2] developed a concept for finding free space on tabletops in order to avoid occlusion by moving objects to free areas. Avoiding occlusion is working fine in settings with much free space available and the objects need to be moved only a bit. However, in highly cluttered environments, where no near-by free space is available, it leads to significant reorganization of the users workspace, which might not be desired. Other research [21] also suggests, that users do not need any reactive occlusion support, as (unwanted) occlusion seldom occurs. Mainly because users arrange their workspace in advance to avoid it or the system can do it by relocation. However, their system did not provide any occlusion support, therefore users might have avoided occlusion by themselves due to the lack of support rather than because they don't do it in general. Also, the setting of a lab bench is rather special and probably not representative for e.g. normal office work. Besides avoiding occlusion, there are concepts to deal with it: Tumble! Splat! by Ramos et al. [15] solve the problem of awareness and access for digital/digital occlusion in drawing applications. Systems like SnapRail [3], occlusion management techniques by Javed et al. [6] or ObjecTop [8] by Khalilbeigi et al. provide means to the user to interact in hybrid settings without losing the awareness of occluded objects but without limiting the user, for example through avoiding occlusion by relocating occluded items. This is done by using interactive proxy objects, e.g. tiny icons representing occluded objects on unoccluded space, allowing for access and perception. However, while helpful for the user, these techniques all consume additional display space on the tabletop. This is not an issue as long as only a low amount of the display is covered, but it might become a problem if there is not much display space left to display the visualization.

3 Initial Study

Study design ProjectTop resolves occlusion using the surface of physical objects as additional display. In order to appropriately top-project digital content on such complex surfaces with various textures, we conducted a user study with 8 participants (2 female, aged 26.5 on average). The goal of this study was to find a representation of occluded objects that conveys as much information as possible while being acceptable for the user. In particular, we wanted to see which properties of the underlying objects are important to the user. We therefore employed three visualization levels, showing different characteristics of the represented object: Low (location only), Mid (location, orientation and type) and High (location, orientation and content).

Apparatus and layout The study was conducted using a LCD-based tabletop with a top-mounted projector. To implement the three visualization levels, we chose a small red dot (the most minimalistic visualization possible) for the “Low”-level, an icon for the “Mid”-level and a miniature version for the “High”-level. These visualizations were projected on a set of physical documents of different types (pure text, text + image, notepad). To see, whether the location of the projection matters, we projected them at two different locations (in-place and moved to the content-free border of the document close to the original position). This leads to a total of 3 visualizations × 2 locations = 6 trials per participant.

Task A set of documents with top-projected content was presented to the participants. In order to force engagement with these documents, the participants had to answer questions related to the documents' content. To find out whether top-projection interferes with interacting with documents, the questions had to be answered in writing on a notepad which also had top-projected content on it. The set of documents was changed for each trial. The study lasted about 20 min per participant. After each trial, the participants rated the visualization regarding the perceived clutter, usefulness, interference between physical and digital content (on a 5 point Likert scale), as well as acceptability of the projection (0 (never) to 10 (continuously)) in the working and storage area.

Results We performed a 2-way ANOVA on the data, Greenhouse-Geisser correction was used in case of violation of sphericity.

The placement of the visualization had significant influence on the perceived clutter ($F_{1,7}$=25, p=.002), the interference (physical: $F_{1,7}$=22.9, p=.002, digital: $F_{1,7}$=17.64, p=.004) and the acceptability of projection in the working area ($F_{1,7}$=21.99, p=.002). There was no significant difference for the storage area ($F_{1,7}$=.13, p=.732) or the helpfulness ($F_{1,7}$=2.03, p=.197). In the significant cases, the users slightly favored the projection at the free spot (marginal means, free-spot/in-place, SE in brackets: clutter: 1.8[.27]/2.6[.19], phys. interference: 1.9[.28]/2.9[.36], dig. int.: 1.7[.27]/2.5[.28], accept. work. area: 7.3[.64]/6.2[.56]). For the acceptability in the storage area, the means were 9.0[.37]/8.96[.35]. While certainly less acceptable, even in working areas, the acceptability of top- projection is more on the acceptable side (> 5.5).

The different visualization had significant influence on the perceived clutter ($F_{2,14}$=19.21, p<.001), the digital interference ($F_{1.2,8.2}$=14.04, p=.001), the helpfulness ($F_{2,14}$=22.53, p<.001) and the acceptability of projection in the working area ($F_{2,14}$=11.47, p=.001). Again, there was no significant difference for the storage area ($F_{2,14}$=2.44, p=0.15) or the physical interference ($F_{2,14}$=.07, p=.84). As expected, the dot visualization introduced minimal clutter (μ=1.4, SE=.175) and digital interference (μ=1.3, SE=.18), however, it was not very helpful (μ=1.75, SE=.31). The icon and miniature visualization introduce more clutter (min.: μ=2.75, SE=.31, icon: μ=2.56, SE=.27), but are perceived as much more helpful (min.: μ=4.25, SE=.21, icon: μ=3.63, SE=.28). All three visualizations received values above 5.5 for the acceptability in the working area and above 8.7 for the storage area, leading to the conclusion that they are all acceptable for most users.

Based on these results, we developed a set of visualization and interaction techniques for hybrid occlusion resolution using top-projection named ProjecTop.

4 ProjecTop

The underlying idea of ProjecTop is to exploit the physical occluder's surface to extend the display and interaction spaces of tabletops in a hybrid physical-digital setting. To do so, we augment back-projected or active LCD based tabletops with projection and depth camera facilities that allow projecting digital content on the surface of physical objects (as peripheral display) and interaction with these projections using direct multitouch gestures.

Fig. 2. Lamp enhanced tabletop

While we use a rather complex setup for our implementation (see figure 9), the advances in miniaturization of sensors and projectors will enable more compact and ergonomic implementations in the future (for instance as an integrated desk-lamp (c.f. figure 2, similar to LuminAR [12]). In the following, we describe visualization and interaction techniques supported in ProjecTop.

4.1 Representing Digital Objects

Based on the study results, we opted for the icon as the default representation. The reason behind this is that the icon has been shown to be acceptable and provides sufficient level of detail about occluded items while introducing less visual clutter to the workspace [6,8]. For ProjecTop, we further extend the icon representation from the study by adding an outline visualization showing the occluded object's outline (see figure 3). The reason for adding the object's outline is that thereby, even more information can be conveyed: Not only the objects position, orientation and type are visible, but also the object's size and possible z-layering are immediately visible to the user. This allows the user to quickly see whether a relevant object - for example a long e-mail - is under a physical object without having to take any physical action (see figure 3(a)). Also, previous research [8] shows that users tend to prefer moving lightweight objects instead of using the provided system support. The spatial awareness (position, size and z-layering are visible) provided by this visualization gives them a hint where the occluded objects are, allowing a more directed and efficient movement of the occluder. Another argument for using the outline element in the visualization is that it allows to “complete” objects in partial occlusion situations (e.g only half of the object is underneath a physical one) as shown in figure 3(b).

The ProjecTop system starts to visualize the occlusion when a meaningful part of a digital object is underneath a physical object. This threshold ensures that digital objects which are, for instance, only with the tip of a corner under an physical object, do not cause disturbing projection artefacts.

Physical objects do not always occlude digital content fully, but partially. This might even be a wanted case of occlusion, as only a part of the occluded object is actually needed, e.g. a small figure on a large page. In such situations, the ProjecTop approach has several advantages over tabletop-based approaches [6,3,8] which use proxy objects like halos or icons on the tabletop display. These approaches require to define a point where the transition between showing the actual object and showing the

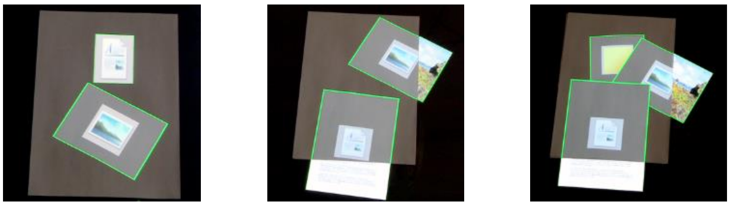

(a) direct overview (b) partial occlusion (c) overlapping

Fig. 3. Visualization overview

icon representation happens (e.g. if a digital object is occluded by 50%, the icon is shown instead of the object). However, it is not clear, when the user actually wants this switching to happen: As mentioned, it might be useful and wanted to have only a part (e.g. 10%) of a technical figure looking out underneath a physical object, while for other objects, an earlier transition is desired (e.g. if 20% are occluded). Additionally, when the switching occurs, there is a notable visual break when the large object changes to a tiny icon. In contrast, ProjecTop allows a continuous transition between fully occluded and fully visible without any distracting abrupt change in the visualization as both styles can coexist in a meaningful way, as the outline can complement the still visible part.

In case of multiple overlapping objects being occluded, ProjecTop visualization conveys their z-order in a natural way, as the outlines occlude each other just like the represented objects do. The aforementioned smooth transition helps to identify the z-order even for two partially occluded items if the overlap happens under a physical object as shown in figure 3(c).

The visualization is not only a static means to provide awareness about hidden objects, but also an interactive means to enable occlusion resolution without having to interact with the physical occluder. Besides retrieving the digital object from underneath the physical one, ProjecTop provides a set of multitouch gestures to interact with the occluded digital objects, even though the occluder. These techniques are explained in the following section.

4.2 Interacting with Occluded Objects

In order to provide a consistent interface, ProjecTop allows standard object manipulation (moving, rotating, etc.) not only on the tabletop's display but also on the surface of physical objects. This allows for instance to naturally manipulate an occluded object "through" the occluding object. A short single finger tap gesture allows to change the z-order of digital objects by making the tapped one the foremost. This gesture is also supported on the tabletop surface as well as on the physical objects. To support immediate access to and fluent organization of digital objects in hybrid settings, we designed three techniques described in the following:

Hold-To-Peek is a quick way of glimpsing at occluded objects. When selecting and holding an occluded object with one finger for a short amount of time, the occluding physical object becomes "transparent" by filling the outline representation with the underlying object's actual content. This allows users to see through the occluder and look at the digital item. When lifting the finger, the object returns to its default outline representation (see figure 4 for an example) This technique can be further extended by replacing the purely visual fade-in with a semantic fading, where gradually more information gets revealed as the user holds longer.

This approach has several advantages over the prior approaches that are based on pulling the icons out from underneath the occluder: Firstly, tapping and holding is a very fast gesture, as it does not require any further movement. Secondly, the gesture can be performed either on the occluder or on the tabletop, which mitigates the problem of the tabletop surface being inaccessible (e.g. between two tall objects). As a result of showing the occluded object in-place, there is also no additional free display space needed next to the occluder, which is beneficial in very cluttered settings.

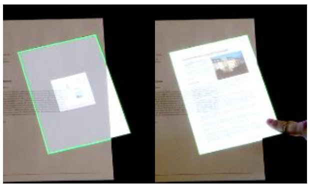

Fig. 4. Hold-To-Peek

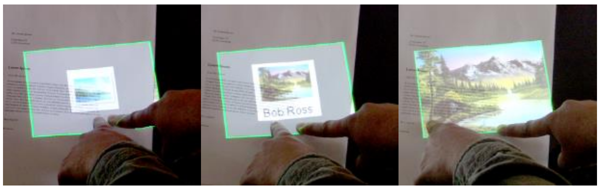

Fig. 5. Semantic-Pinch

Fig. 6. Direct HyperMove

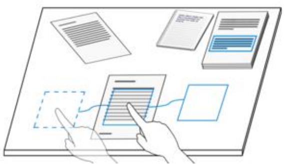

Fig. 7. Manual Hypermove

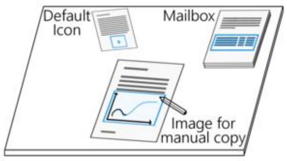

Fig. 8. Semantic-Pinch use

SemanticPinch provides a permanent access to higher levels of detail about the occluded object. Similar to the typical pinch-to-zoom gesture, SemanticPinch is used to change the representation of an occluded digital object on a physical one (see figure 5): As the user "zooms in", more and more detail about the object is revealed. Contrary to the state-of-the-art semantic zooming facilities for occlusion [8], the visualization remains in its current semantic zoom representation when the user lifts the fingers. This way, the user is able to control the amount of information he wants to receive about an object permanently.

Using this technique, users can maintain a more detailed view on occluded objects. Practical uses-cases are for example an occluded mailbox which the user can semantically zoom in to view the last n received mails instead of just a mailbox icon. Due to the top-projection, he has an always-on-top mailbox overview. Another use-case would be tracing digital objects to copy them on physical paper: The user can place the digital object, occlude it with a sheet of paper and then semantically zoom in until the full object is shown, allowing him to trace it using a pen. See figure 8 for an illustration.

Moving objects is a key action for workspace organization and well supported on digital tabletops. With the addition of physical objects, moving digital objects becomes cumbersome: The direct moving path might be obstructed by a physical object, requiring the user to either move the physical object or drag the digital one around the physical one. In order to better support these cases, we developed HyperMove.

HyperMove is an extension of the moving gesture supported on conventional multitouch surfaces and is not bound to the tabletops display, but extends the interaction space onto the surface of physical objects. ProjecTop supports two variants of HyperMove: manual and direct. The manual version extends the conventional dragging to be used across physical objects: The user can for instance start dragging a digital object on the screen and move it across a printed page while searching for a suitable place. This variant is particularly useful if the user does not know the exact place where he wants the object and if there are only flat objects in the path.

The direct version uses a bi-manual gesture to move digital items over large distances without having to drag them: The user references the desired position of the object by a static tripod-gesture (similar to teleport in [8]) with one hand and can then move digital objects immediately to this location by performing a single finger flick gesture with the other hand on them. Again, ProjecTop allows this gesture to be performed on the tabletop as well as physical objects (see figure 6 and 7). Besides supporting conventional moving tasks, this gesture allows for more advanced an quick organization, like piling several digital objects under a stack of physical ones.

5 Implementation

All techniques described above are implemented in a system C# using DirectX. It runs on a Samsung SUR40 tabletop system with a Full HD screen resolution of 1920x1080. The raw infrared image from PixelSense is used to do a shadow tracking of physical objects on the table's surface. In addition, fiducial tag markers are used for identifying physical objects. ProjecTop uses a standard Full HD projector mounted above the tabletop to realize the top-projection. To align the projection with the tabletop coordinate system and thereby allow precise projection on the physical objects, a one-time calibration is needed if the table or projector are moved. The touch input is provided by the SUR40 for the tabletop surface, and by a depth camera (Microsoft Kinect) based touch recognition framework for the passive physical objects. Our frameworks approach is an extended version of the approach used by for dSensingNI [9]. We added the ability to allow object and touch recognition in areas where the background image contains unknown values and simple background subtraction does not work.

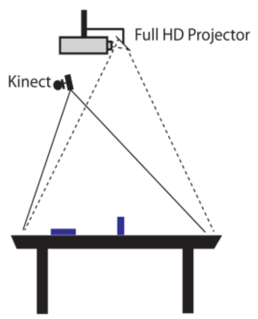

Fig. 9. Lab setup for prototype

6 User Study

The overarching goal of the study was to find out how ProjecTop's features, i.e. the in-place resolution and the visibility of the spatial layout perform under different clutter levels. Performance in this context mainly refers to the time that users need to find an occluded object and the effort required for doing so (e.g. the number of interactions with physical objects).

Based on prior studies [6,8], we decided to compare the top-projection based ProjecTop approach with a system named “Baseline”, that relies purely on the tabletop's display. Baseline consists of two well-established occlusion support techniques: occluded objects are visualized as icons around the occluder to provide accessibility and a glowing [6] visualization surrounding the occluding object to provide general awareness. This way we can ensure generalizability and comparability of our study's results. Additionally, we did not impose any restrictions on the interaction with the physical objects to force usage of the digital occlusion resolution techniques. The users are allowed to lift or move physical objects if they prefer to do so over using the provided system support in both systems throughout the study.

As we know from previous work, the usage and performance of occlusion related techniques is dependent on the amount of clutter of the tabletop [8]. To see whether this has any impact in a top-projection scenario, we had two conditions (low and high clutter) in our evaluation, differing in the number of physical and digital objects present on the tabletop. We rank the clutter level of hybrid tabletop settings as a function of the number of digital objects on a tabletop and the amount of occluded display surface. In this experiment, for the low clutter conditions, we used 10 digital objects (of which at least 5 were initially occluded) and 3 physical objects. This choice lead to about 40% occluded display area. The high clutter condition consisted of 15 digital objects (10 were occluded) and 5 physical objects leading to about 60% occlusion.

6.1 Apparatus and Material

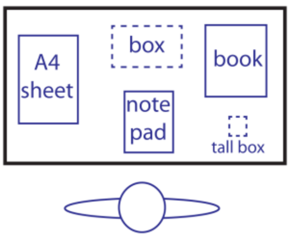

Fig. 10. Initial setup. The dashed objects were only present in high clutter conditions.

We conducted the study using the lab hardware setup described in the implementation section. As physical objects, we used the following items: A thick book (20 by 24 cm with 4 cm height), a standard DIN A5 paper notepad, a printed A4 document, a cardboard box sized approximately A5 with a height of 3 cm and a heavy small high box (7 by 7 cm with a height of 14 cm). To have different surface textures for the projection, the document was printed with a lighter texture on it, the book was chosen to have a darker one. The small box and the cardboard box were only present in the high clutter conditions. The other objects were present in both, high and low clutter, conditions.

We chose the objects to be representative for an office desk with respect to size and ease of moving them. As the size of the tabletop is limited and the display could be dam aged when high loads are placed on it, we had to abstain from other common objects like monitors, or a full keyboard. The initial physical layout for the study can be seen in figure 10.

As for digital objects, we chose a set of photos, documents and virtual Post-it notes. For each of these type, different categories existed, e.g. for the documents, there were documents about persons, movies, universities, etc. The objects were randomly chosen from the pool and placed randomly across the tabletop surface at the beginning of each trial.

6.2 Task

The experiment consisted of a search task in which the participants had to find and count a set of digital target objects, for example the number of documents about movies. The targets were mixed with other digital objects of different types or topics (distractors). All target objects were initially hidden under physical objects.

6.3 Hypothesis

Compared to Baseline, we expected that ProjecTop reduces the time needed to find a specific object as, in most situations, the occlusion can be directly resolved without having to perform a move gesture (e.g. for a semantic zoom) but rather a single tap-and-hold. Due to the in-place occlusion resolution, we assumed that the users are much more unlikely to actually drag out objects from underneath the occluder, as inspecting the full object is directly possible. Additionally, the interaction with physical objects should be reduced as ProjecTop provides an easy way to see where an occluded object is placed before lifting or moving the physical occluder. Therefore these actions can be carried out more efficiently. In sum, the hypotheses to be examined were:

- H1: ProjecTop requires less time to find and access objects
- H2: ProjecTop reduces the number of full object pull-outs (the occluded object is fully moved to unoccluded space)
- H3: ProjecTop reduces the number of object interactions and the distance objects are moved

6.4 Experiment Design

The experiment was designed as a two interface (Baseline and ProjecTop) × two clutter levels (Low and High) within subjects experiment. The presentation order for the four conditions was counterbalanced. For each condition, there were five repetitions with different types (documents about movies, images of persons, …) and numbers (one to three) of targets. For each condition, the order of the target types for the five trials was randomly assigned for each participant. Prior to the each set of trials, the users got a short introduction on the current interface and task and had some time to play around with the system until they felt comfortable using it. The users where then asked to perform the counting task as accurate and fast as possible.

The main measurement for the task was the completion time that was measured by pressing the space bar of a keyboard within the participants' arm reach when starting and after finishing the counting. The number of interactions with physical objects as well as their movement distance was logged. For the digital objects the number and type of special interactions, like pulling out from underneath an occluder, was logged. The number of targets determined by the participants was recorded to determine the error-rate made by the participants. The participants advanced to the next trial without having been told if they were right, in order to not influence their behavior (e.g. becoming faster or slower).

After each condition, the participants had to fill in a NASA-TLX [14] questionnaire to measure the perceived workload. At the end of each session, the participants were interviewed in a semi-structured way to gather additional subjective feedback. We put an emphasis on the appropriateness of the projection during the interview. The whole sessions were recorded on video and lasted about one hour.

6.5 Participants

17 participants (15 male, 2 female, all but 2 right-handed) aged between 21 and 46 (μ=27.588, σ=5.657) were recruited from the local university. There was no compensation provided. All participants had previous experience with touch interfaces and all except three participants had previous experience with interactive tabletops.

There was a total of 17 participants × 2 interface × 2 clutter levels × 5 repetitions = 340 trials

7 Results and Discussion

After preprocessing of the recorded trial data (e.g. removing outliers more than three standard deviations away from the mean), 9 trials had to be removed due to errors in time measurement. To analyze the data of the remaining 331 trials, we used a 2-way repeated measures ANOVA. There was no significant influence of the presentation order of the tasks for any measure, e.g. any learning effect on the interface. Therefore, the presentation order was neglected in the following analysis. For all pair-wise comparisons, Bonferroni correction was applied where necessary. The Shapiro-Wilk test was used to check for normal distribution and ensure applicability of ANOVA.

7.1 Task duration and errors

The analysis of the measured trial time revealed that the interface as well as the clutter level had significant influence. ProjecTop (μ=16.5s, SE=1.7) performed significantly

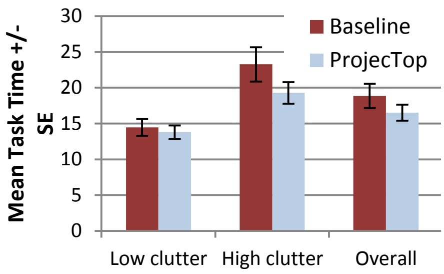

Fig. 11. Average trial duration by interface

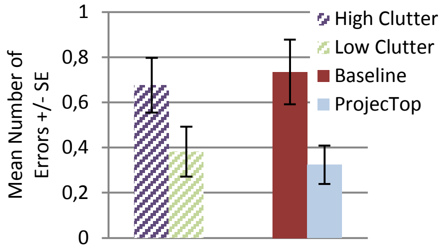

Fig. 12. Average number of counting errors per trial by clutter level (left) and by interface (right)

(F1,16=6.09, p=0.025) faster than Baseline (μ=18.9s, SE=1.1), see figure 11. The clutter level also had a significant (F1,16=34.52, p<0.001) influence on the trial time (μ=14.1s, SE=1 for low clutter, μ=21.3s, SE=1.9 for high clutter). Additionally, there was a significant effect of interface × clutter on the completion time (F1,16=6.67, p=0.02). Pairwise comparison revealed, that ProjecTop performed significantly faster under the high clutter condition (p=0.017). For the low-clutter condition, the difference was not significant (p=0.29).

Interestingly, when looking at the average number of wrong answers given by the participants, there was also a significant ($F_{1,16}$=8.209, p=0.011) difference between ProjecTop and Baseline in favor of ProjecTop (see figure 12 right). The influence of the clutter level was also significant ($F_{1,16}$=4.68, p=0.046) showing a higher number of errors for high clutter conditions as expected (see figure 12 left). There was no significant effect of interface × clutter on the number of errors.

These results show that ProjecTop performs faster than Baseline, especially in highly cluttered environments. This supports H1. Interestingly, despite being faster, the participants made significantly less mistakes (reporting a wrong target count) using ProjecTop than they did with Baseline. One reason for this might be that with Baseline, icons which were shadowed by a thick or high object were not visible from the users point of view. Therefore; there was a higher chance that those are missed (leading to more errors) or users needed to take special care of them by explicitly looking behind objects (leading to additional time). ProjecTop inherently circumvents this problem as the visualization is right on the occluder rather than behind it.

7.2 Interactions with Objects

When analyzing the number of interactions with physical objects, it turned out, that the interface had a significant influence on the number of interactions ($F_{1,16}$=9.668, p=0.007) with physical objects and the distance physical objects were moved ($F_{1,16}$=5.72, p=0.029). In ProjecTop conditions, the number of interactions was slightly higher with an average of 1.38 interactions per object, whereas Baseline had an average of 1.07. The same holds for the distance objects were moved: Baseline conditions had an average distance of 62.99 mm per object - ProjecTop conditions had

96.55 mm. Hence, H3 is rejected. There was no significant effect of the clutter level or interface × clutter.

Interestingly, even though the amount and distance objects were moved was higher with ProjecTop, it was less time consuming. A possible explanation is the spatial awareness provided by ProjecTop's visualization. It gives the users a strong visual clue of the spatial layout of the occluded digital objects before interacting (e.g. moving or lifting) with the physical occluder. As a result, no additional time is needed to interpret the layout underneath the occluder.

With respect to H2; it turned out that the number of full object pull-outs was a bit higher with ProjecTop (10.62 for Baseline and 11.62 with ProjecTop). However, with $F_{1,16}$=0.73 and a resulting significance of p=0.407 we assume that the interface had no real influence on the number of objects pulled out. Analysis of the recorded video data revealed that most users tended to lift most of the physical occluders before they finish, thereby "resolving" the occlusion for all underlying digital objects. Therefore, H2 is rejected.

7.3 Task Load Index

The analysis of the NASA TLX Questionnaire revealed that the interface had no significant influence on the measured categories (see figure 13 for an overview). There was also no significant effect of interface × clutter except for the perceived effort ($F_{1,16}$=4.848, p=0.043): Baseline has values of 9.82 for the high clutter condition and 6.18 for low (p=0.011). In contrast, ProjecTop has 7.71 for high and 7.12 for low clutter respectively (p=0.365). The clutter level had significant influence on all measures except performance (p=0.17), with mental demand (p=0.005), physical demand (p=0.037), temporal demand (p=0.012), effort (p=0.01) and frustration (p=0.055) being lower for the low clutter conditions as one would expect.

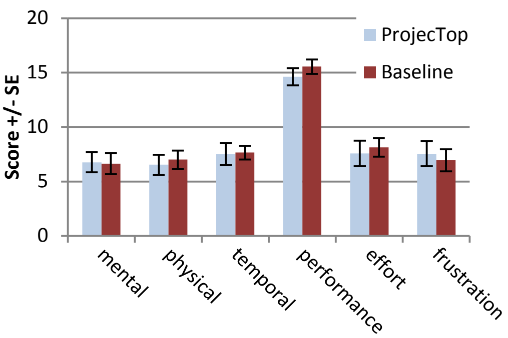

Fig. 13. Average trial duration by interface

7.4 User Feedback

The participants' feedback in the post-interviews regarding top-projection varied: Three users found the ProjecTop visualization to be *"chaotic"* (P11) compared to the baseline. A few participants explicitly criticized the interference between projected digital and physical text on the objects, which both get unreadable if they overlap. However, most participants really liked top-projection *"ProjecTop's projection on objects is very cool"* (P10), some even wanted the full content of digital objects to be always projected without user interaction (P10, P15, P16, P17), instead of the icons. P12 found ProjecTop to be *"clearer [than baseline] with many objects"* and others stated that they had difficulties to obtain an overview in the Baseline condition (e.g. P15: *"Did I already take a look at that icon or not?"*).

Top-projection A common statement was that top-projected visualization should only be used for the storage area, but not for the working area (e.g. P12: “Projection is only good in storage areas and not for the working area, because one concentrates there on specific objects and doesn't want to overlook the mass of objects”). At the same time, the concern was raised that the question of where to draw the border between working and storage area is di-cult to answer. To resolve this problem, users suggested a differentiation by focus, so that physical objects the user interacts with do not get top-projection, while the others do (P2, P10). An alternative suggestion was to vary the brightness of top-projection, instead of deactivating it, so that for objects in focus, the projection gets darker to improve readability of the physical content (P15). P16 suggested a combined Baseline/ProjecTop system where the system decides whether objects are suitable for top-projection (then using ProjecTop) or not (using Baseline style). The user should then, however, be able to switch between the styles for specific objects manually. In the context of a combined system, P2 wanted to have the Hold-To-Peek gesture for the icons in Baseline, so the actual object gets top-projected when touching its icon. Overall, it seems beneficial to combine both approaches.

Fig. 14. Differently colored outlines for categorization

Accessing occluded objects All users found accessing occluded object through ProjecTop easy. The Hold-To-Peek gesture was even found to be sufficient to complete the task by some users, e.g. *“ProjecTop was often sufficient to see what's underneath”* (P6). However, there occurred problems in case of digital-digital occlusion, which is currently not resolved by ProjecTop. In case of one digital object being fully occluded by another under a physical object, the occlusion was reflected by ProjecTop's visualization, making the bottom digital object inaccessible. Baseline's icon representation inherently resolves this occlusion, as the icons do not overlap. However, ProjecTop could be extended to provide access in such cases, too. Additionally, P14 suggested to add different outline colors to group objects into categories and further speed-up finding objects as shown in figure 14.

Spatial awareness In general, the participants found ProjecTop's spatial awareness useful, especially if they wanted to lift the occluder to access one or more occluded objects: *“Spatial awareness was very helpful, I could look at the physical object and directly saw where the objects are, with Baseline I had to look around the [physical] object to see where [digital] objects are.”* (P15). Besides the concerns regarding interference, there were no negative comments on it. Touch on objects Touch on arbitrary physical objects was strange to two users: P4 found *“touching flat objects more intuitive than thick objects, as flat one are closer to the screen which is naturally interacted using touch. For thick or uneven objects, it was strange to touch the surface in order to interact.”* P1 said that, *“for thick objects, there is a problem if a gesture leaves the object, for instance when dragging out an object from underneath a book, as the touch surface is not continuous at the books edge.”*. The other participants were positive about it, so it can be assumed to be an appropriate way of interacting with top-projected content.

8 Implications

Based on the study results, we propose three main implications for informing the design hybrid tabletops augmented with top-projection:

How to resolve occlusion? Top-projection and in-situ or on the tabletop display and off-place? The study showed that, in general, top-projection is a suitable means to extend the display and interaction spaces in hybrid settings and particularly to resolve physical/digital occlusion. However, its use has to be considered carefully to support the user without distracting him or hindering his work. As a result, it cannot be applied to every object to avoid conflicts with the user's task and needs. One should therefore decide on a per-object basis whether to use top-projection or not. This decision is influenced by two main factors: The user's current task and preferences on one side and the occluders properties on the other.

The first factor is driven by the user’s needs during a task, e.g. top-projection on a document the user currently reads is clearly unwanted. However, reading and actively working with objects mainly happens in the working area of the table. One can use the knowledge about the users working area to automatically switch to top-projection when an object is within this area, only using top-projection within the storage areas, as P12's comment suggested.

The second factor is due to the fact that objects are not suitable for projection and/or interaction. This can either be because they very dark or strongly textured, making the projection unrecognizable, or because the objects use touch as input modality (e.g. keyboard or tablet), which would interfere with the user wanting to interact with the projection. Again, switching to a tabletop-based representation is a possible solution. The system can be enabled to recognize common unsuitable objects and do the switching automatically.

Besides these two factors which can be considered programmatically, it is sensible to provide the user a means of individually controlling top projection, so he can turn it on or o for specific objects at will. Also, some users said they would prefer to have a gesture to globally toggle top-projection. Then, they can use top-projection like an expose mode when the need an overview without being distracted otherwise.

Is spatial awareness useful? Despite the higher movement count of physical objects and a large number of interactions, ProjecTop was significantly faster in the high clutter condition. Video analysis confirmed that users often lifted all objects in Baseline and ProjecTop conditions prior to finishing the task in order to check for missed objects. Probably, even if unconscious, the spatial awareness in ProjecTop facilitates this process as participants already knew the underlying layout, leading to the faster completion times and less errors. It is therefore important to provide such visual clues, even if users don't consciously use them and keep interacting with the physical objects to resolve occlusion.

How to deal with touch across tall objects? A problem with expanding the interaction space onto the physical objects is the incontinuity of the touch surface, two participants brought up: The steep edges of physical objects make smooth touch interactions impossible, as the user has to lift his finger in a step like manner to always be on the touch surface. A possible solution to mitigate this problem can be a kind of hover-touch, which allows to continue the move while having the finger in the air. This enables a smoother transition between the surfaces, as the user can start lifting the finger before he reaches the object's edge.

9 Conclusion

Solving the problems hybrid tabletop settings pose to interaction and UI design is one key aspect to the success of tabletops in everyday environments. For instance in offices, interactive table would not only be used on its own when integrated with the office's furniture and therefore being used like a conventional desk. Physical objects which are placed on the table are - unlike objects in the context of tangible user interfaces, which are part of the interface itself – not part of the actual user interface, making them a foreign object to the interface.

We presented ProjecTop, an in-place physical/digital occlusion resolution system. It does not only circumvent occlusion by using the tabletop's display, but also uses the physical objects surface in order to do so. The formerly foreign object is made part of the interface by using it as an additional interactive display surface to resolve the occlusion situation in place, therefore not consuming additional display space on the high-resolution tabletop display. We further reported an evaluation of the ProjecTop concept, showing that the top-projection approach is not only beneficial for the resolution of occlusion related problems like awareness and accessing, but also to be a valuable addition to interactive surfaces if used advisedly. As the qualitative feedback showed, the spatial awareness is perceived as useful by the participants. Future studies should be conducted to further explore this aspect.

10 Acknowledgements

This work was partially supported by the ICT R&D program of MSIP/IITP. [13-921-03-001, Development of Smart Space to promote the Immersive Screen Media Service]

11 References

1. Cotting, D., Gross, M.: Interactive environment-aware display bubbles. In: Proc. UIST. pp. 245-254. ACM, New York (2006)
2. Freeman, E., Brewster, S.: Messy tabletops: Clearing up the occlusion problem. In: CHI EA. pp. 1515-1520. ACM, New York (2013)
3. Furumi, G., Sakamoto, D., Igarashi, T.: Snaprail: a tabletop user interface widget for addressing occlusion by physical objects. In: Proc. ITS. pp. 193-196. ACM, New York (2012)
4. Hartmann, B., Morris, M.R., Benko, H., Wilson, A.D.: Pictionaire: Supporting collaborative design work by integrating physical and digital artifacts. In: Proc. CSCW. pp. 421-424. ACM, New York (2010)

5. Iwai, D., Sato, K.: Limpid desk: see-through access to disorderly desktop in projection-based mixed reality. In: Proc. VRST. pp. 112-115. ACM, New York (2006)
6. Javed, W., Kim, K., Ghani, S., Elmqvist, N.: Evaluating physical/virtual occlusion management techniques for horizontal displays. In: P. Campos et al. (Eds.) INTERACT, LNCS, vol. 6948, pp. 391-408. Springer, Heidelberg (2011)
7. Khalilbeigi, M., Steimle, J., Mühlhäuser, M.: Interaction techniques for hybrid piles of documents on interactive tabletops. In: CHI EA. pp. 3943-3948. ACM, New York (2010)
8. Khalilbeigi, M., Steimle, J., Riemann, J., Dezfuli, N., Mühlhäuser, M., Hollan, J.D.: Objectop: occlusion awareness of physical objects on interactive tabletops. In: Proc. ITS. pp. 255-264. ACM, New York (2013)
9. Klompmaker, F., Nebe, K., Fast, A.: dSensingNI: A framework for advanced tangible interaction using a depth camera. In: Proc. TEI. pp. 217-224, ACM, New York (2012)
10. Koike, H., Sato, Y., Kobayashi, Y.: Integrating paper and digital information on enhanceddesk: a method for realtime finger tracking on an augmented desk system. ACM TOCHI 8(4), 307-322. ACM, New York (2001)
11. Liao, C., Tang, H., Liu, Q., Chiu, P., Chen, F.: Fact: Fine-grained cross-media interaction with documents via a portable hybrid paper-laptop interface. In: Proc. ACM MM. pp. 361-370. ACM, New York (2010)
12. Linder, N., Maes, P.: Luminar: Portable robotic augmented reality interface design and prototype. In: Adj. Proc. UIST. pp. 395-396. ACM, New York (2010)
13. Mitsuhara, H., Yano, Y., Moriyama, T.: Paper-top interface for supporting notetaking and its preliminary experiment. In: IEEE SMC. pp. 3456-3462 (2010)
14. NASA Ames Research Center, Moffet Field: Nasa tlx (1988)
15. Ramos, G., Robertson, G., Czerwinski, M., Tan, D., Baudisch, P., Hinckley, K., Agrawala, M.: Tumble! splat! helping users access and manipulate occluded content in 2d drawings. In: Proc. ACM AVI. pp. 428-435. ACM, New York (2006)
16. Rekimoto, J., Saitoh, M.: Augmented surfaces: A spatially continuous work space for hybrid computing environments. In: Proc. CHI. pp. 378-385, ACM, New York (1999)
17. Spindler, M.: Spatially aware tangible display interaction in a tabletop environment. In: Proc. ITS. pp. 277-282. ACM, New York (2012)
18. Spindler, M., Martsch, M., Dachselt, R.: Going beyond the surface: Studying multilayer interaction above the tabletop. In: Proc. CHI. pp. 1277-1286. ACM, New York (2012)
19. Spindler, M., Stellmach, S., Dachselt, R.: Paperlens. In: Proc. ITS. pp. 69-76. ACM, New York (2009)
20. Steimle, J., Khalilbeigi, M., Mühlhäuser, M., Hollan, J.D.: Physical and digital media usage patterns on interactive tabletop surfaces. In: Proc. ITS. pp. 167-176. ACM, New York (2010)
21. Tabard, A., Gurn, S., Butz, A., Bardram, J.: A case study of object and occlusion management on the elabbench, a mixed physical/digital tabletop. In: Proc. ITS. pp. 251-254. ACM, New York (2013)
22. Wellner, P.: Interacting with paper on the digitaldesk. Com. ACM pp. 87-96. ACM, New York (1993)