

Is Two Enough?! Studying Benefits, Barriers, and Biases of Multi-Tablet Use for Collaborative Visualization

Thomas Plank¹, Hans-Christian Jetter¹, Roman Rädle², Clemens Nylandsted Klokmose², Thomas Luger¹, and Harald Reiterer³

¹Univ. of Applied Sciences Upper Austria {thomas.plank,hans-christian.jetter,thomas.luger}@fh-ooe.at

²Aarhus University, roman.raedle@cc.au.dk, clemens@cavi.au.dk

³Human-Computer Interaction Group, University of Konstanz, harald.reiterer@uni-konstanz.de




Figure 1. Two participants of our study working on a study task (A). Different visualizations of OSCE data: Map (B), Parallel Coordinates (C), Data Table (D).

ABSTRACT

A sizable part of HCI research on cross-device interaction is driven by the vision of users conducting complex knowledge work seamlessly across multiple mobile devices. This is based on the Weiserian assumption that people will be inclined to distribute their work across multiple “pads” if such are available. We observed that this is not the reality today, even when devices were in abundance. We present a study with 24 participants in 12 dyads completing a collaborative visualization task with up to six tablets. They could choose between three different visualization types to answer questions about economic data. Tasks were designed to afford simultaneous use of tablets, either with linked or independent views. We found that users typically utilized only one tablet per user. A quantitative and qualitative analysis revealed a “legacy bias” that introduced barriers for using more tablets and reduced the overall benefit of multi-device visualization.

Author Keywords

cross-device interaction; group work; information visualization; multiple coordinated views; tablets

ACM Classification Keywords

H.5.3 Information Interfaces and Presentation (e.g. HCI): Collaborative Computing

INTRODUCTION

A sizable part of HCI research on cross-device interaction is driven by the vision of enabling users to conduct complex knowledge work (e.g., active reading [9, 10], brainstorming [30], sensemaking [16]) by seamlessly working across multiple mobile devices. In this vision, multiple tablets or phones form seamless and powerful device “communities” [39] or “symphonies” [16] that can enhance or even replace traditional desktop devices or large screens. This echoes Weiser’s vision of ubiquitous computing where mobile “pad” devices (i.e. tablets) are a commodity similar to sheets of paper and are readily available in our workplaces with up to ten or twenty pads per room [54]. Thereby Weiser’s pads are not necessarily owned by individuals but are readily available to be used, passed around, and shared whenever needed.

Today, however, our tablets play a quite different role and we usually do not have 10-20 tablets per room. They are predominantly personal devices for solitary use and not Weiserian pads that are easily shared between users or spread out like documents on a desk. Instead, they are used to access personal information such as social media or photos [3] and tech companies sell them at a high unit cost without great interest in letting users share them [4]. Using multiple tablets at a time or sharing them with others therefore rarely happens outside research labs. This has important implications for the design of multi-tablet systems: Even when plenty of inexpensive tablets will be available in future, users could be strongly biased towards solitary use. Users might need a long time to get used to the idea of sharing and using multiple tablets in parallel and also may need help through the interface to fully understand and exploit their benefits.

We therefore conducted a study to observe if users would benefit from multiple tablets during collaborative visualization without prior training or experience in multi-tablet use.

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 ACM 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 2017, May 06–11, 2017, Denver, CO, USA.

© 2017 ACM. ISBN 978-1-4503-4655-9/17/05. . . . \$15.00

DOI: <http://dx.doi.org/10.1145/3025453.3025537>

This question is highly relevant for our research in which we explore better support of collaborative sensemaking with multiple tablets in “bring-your-own-device” and “walk-up-and-use” scenarios. Ideally, the tablets would quickly serve users as physical manifestations of different views on the data and their micro-mobility [31, 33] would make it easier to pass views around, share them with collaborators, organize them in space (e.g. on a desk), and discuss content. However, we also expected barriers and biases like unfamiliar cross-device interactions and working styles that would keep users from recognizing and exploiting the tablets’ full potential.

For our study, we built a multi-tablet data visualization system to study if and how users utilize multiple tablets during a controlled sensemaking task and how differences in working styles and device usage evolve. We observed 12 pairs of users (24 participants) and provided each pair with 6 tablets (Figure 1A). Each tablet could be used to switch between 3 different visualizations (Figure 1B-D) to answer questions about demographic and economic data from European countries. A qualitative and quantitative analysis of user behaviour, task results, task completion times, and interviews were used to identify relevant patterns and to formulate following main findings:

Finding 1: Participants underutilized the available tablets and were hesitant to use multiple tablets in parallel, hinting at a “legacy bias” of using tablets as “computers” rather than “documents”. **Finding 2:** The typical starting configurations were “one tablet per user” and “one tablet per group”. They served as natural entry points and users only temporarily switched to more complex configurations. These configurations introduced a higher mental demand and did not improve efficiency and effectiveness. **Finding 3:** Movements of tablets in space primarily served (i) to improve ergonomics and collaboration and (ii) to establish semantically meaningful orders. However, we did not observe physical piling, filling, or clustering of tablets based on content. **Finding 4:** Participants transferred familiar techniques and tools from working with paper to working with the digital “documents” on a tablet, e.g., using hands or paper as guides or rulers.

This paper presents how we arrived at these findings and discusses their implications for the design of multi-tablet visualization systems. We also contribute to the future analysis of multi-tablet systems by discussing the concept of a *legacy bias* [35] towards single-device use and introducing a *tablet utilization score* as analytical tool.

BACKGROUND & RELATED WORK

Weiser’s idea of ubiquitous computing as “the computer for the 21st century” [54] has been most influential for HCI research. It has inspired scores of researchers who cite Weiser’s vision of technologies that “weave themselves into the fabric of everyday life” [54] as long-term goal of their work. Yet the more we use devices similar to Weiser’s envisioned “tabs”, “pads”, and “boards” instead of desktop personal computers, the more researchers have started to become critical about our interpretation of Weiser’s vision and how we as a research community often fail to question our underlying presumptions [1, 4, 5, 42, 38].

Multi-Device & Cross-Device Interaction

One rarely questioned presumption concerns the role that mobile devices should play in future. Weiser envisioned a future where users “unconsciously” [54] use readily available inch-scale “tabs”, foot-scale “pads”, and yard-scale “boards” and distribute their work across them. “The real power of the concept comes not from any one of these devices; it emerges from the interaction of all of them” [54]. In this ubicomp vision, personal computers (PCs) play the role of a transitional technology. The future of computing is supposed to happen (only) with novel post-PC devices. This belief has become very influential and entrenched in HCI and UbiComp research and a sizable part of publications therefore attempts to realize powerful systems without desktop PCs.

One stream of this research are cross-device systems to support collaborative sensemaking tasks with large interactive walls [2, 22, 43] or interactive tabletops [21, 23, 37, 36, 51]. Some approaches also combine tablets for personal use with tabletops [34, 53, 57]. Recent work, however, has found that users are also able to efficiently collaborate without large shared displays and only with tablet-sized devices [57]. Therefore we are interested in creating such *inexpensive, installation-free, tablet-only* solutions for collaborative visualization. Using multiple mobile devices for such co-located collaboration has already been suggested in the past [29]. For example, Lucero et al. have proposed systems that entirely rely on smartphones for tasks such as brainstorming [30] or photosharing [28] without PCs or laptops. Other researchers propose to join multiple smartphones [6, 47] or tablets [27] to create shared tiled displays. In other work, Hamilton et al. chain complex sensemaking tasks across up to 10 tablets [16] and Rädle et al. and Wozniak et al. propose spatially-aware mobile devices as alternatives to interactive tabletops [39, 56]. Working with multiple co-located tablets is also explored by Hinckley et al. and Marquardt et al. who propose different cross-device interaction techniques that support collaborations [18, 33].

A related stream of research is concerned with understanding existing cross-device practices. For discovering paradigms and issues of new cross-device interaction techniques, Santosa and Wigdor performed a field study and interviewed professionals from a broad range of industries [46]. Jokela et al. analyzed real-life multi-device uses cases based on diaries and interviews and used this data to characterize usage practices of common devices (smartphones, computers, tablets, home media centers), identify main usage patterns and discuss practical challenges of using these devices together. Additionally, they defined three levels of decisions for choosing the device in a particular situation [24]. Despite this existing work, Grubert et al. argue that HCI has just begun to explore this area and that the understanding of the design of multi-device applications and implications from multi-device use in social settings are still “not well researched in the community” [14]. Grubert et al. consider them as fundamental challenges in HCI and prompt for opportunities to further investigate micro-mobility for co-located interactions. We consider our work as a consequential next step into this direction.

Tablets vs. Paper

Weiser anticipated that his “pads” will “[...] behave something like a sheet of paper (or a book or a magazine) [...]” [54] and will be used in a similar fashion. Product names like “iPad” also reflect the assumed similarity between using tablets and using a pad of paper. This presumption is rightfully challenged by Haber et al. who conducted a study of the differences and commonalities between using tablets and paper [15]. They found that paper is still preferred and that introducing tablets affects the levels of verbal communication and gaze engagement. They call for a better understanding of the subtleties entailed when tablets are used instead of paper.

Chen et al. studied the month-long deployment of a multi-slate e-paper system for active reading [9, 10]. They list some advantages of paper over their digital system [9]. For example, digital documents have limited annotation & highlighting functionality, are less portable, need batteries, and lack the physical qualities of paper that support quickly skimming through lengthy documents. During the deployment of their system, Chen et al. also observed that no users carried around more than two of the four slates [10]. For future work, they proposed to further explore the use of multiple interlinked mobile displays in other domains, for example for the promising area of information visualization given its use of multiple linked data views [10]. Our study was strongly inspired by this proposal and we consider our study as the first step towards an in-the-wild deployment of a multi-tablet visualization system.

Tablets for Visualization

Büring et al. proposed using scatter plots with different distortion techniques for visualizing large data sets on small screens, e.g., PDAs [8]. In subsequent work, Sadana et al. designed and implemented multi-touch scatter plots for tablets [44]. Later they also proposed a design for tablet-based multiple coordinated visualizations [45]. Their design is based on maximizing screen space for multiple views on a single tablet but, unlike our system, does not involve further tablets for cross-device multiple views. This is also true for TouchViz by Drucker et al. [12], a tablet-based data analytics system that uses touch actions to eliminate traditional control panels to save screen space.

Inspired by this variety of previous and related work, we decided to contribute to a better understanding of multi-tablet use during co-located collaboration by investigating how users utilize multiple tablets during a collaborative information visualization task. In the following section, we describe the design of our study and prototype in detail.

STUDY

For the study, we designed and implemented a prototype system for visually exploring a data set with multiple tablets. The web-based system used JavaScript frameworks such as D3.js [7] and Meteor to visualize yearly key figures from socio-economic data (e.g. fertility rate, unemployment rate) for 24 European countries provided by the OSCE for 2003 to 2011. This data was presented using three visualizations:

Visualization 1 was a zoomable geographic map that enabled participants to show country details and map different key figures to a country’s color. The color mapping helped to identify geographic patterns or outliers (Figure 1B). *Visualization 2* was a parallel coordinates plot of country data. The participants could rearrange or remove axes and highlight single countries. The plot helped to identify correlations between key figures within and across countries (Figure 1C). *Visualization 3* was a table with countries in columns and the different key figures for different years in rows. To select and highlight a country, users could tap a column in the table. The table enabled users to look up and compare specific key figures for specific years (Figure 1D). All visualizations could be filtered with dynamic queries to display figures only for a single year or to show averages for a range of years.

We provided pairs of users with 6 tablets and used a *within-subjects design* to compare two conditions: SINGLE and PAIRED. In PAIRED mode, each of the 6 tablets was linked to a counterpart forming 3 pairs of 2 tablets. For each pair the selected country and year (or range of years) were kept in sync while the visualization type and parameters could be controlled individually. Both devices of each pair were visibly marked with a shared symbol (triangle, square, circle) on a sticker to let participants see which tablets were linked. In SINGLE mode, there was no linking and all tablets provided entirely independent views. Following the taxonomy of Terrenghi et al. for multi-person-display ecosystems [52], our study used a foot size ecosystem of 6 tablets with one-one social interaction for which the displays were coupled in PAIRED mode and uncoupled in SINGLE mode.

Tasks

We designed tasks that were best solved by simultaneously using multiple tablets for parallel views of the data. Typically, participants had to compare figures across countries or identify temporal trends and geographic patterns (e.g. north-south divide). For instance, the participants had to answer the following question about data from 2007: “In which country west of France was the native-born unemployment rate higher than the total unemployment rate?”. To solve this efficiently, participants could combine multiple tablets with maps to identify the countries west of France (Spain, Portugal, UK, Ireland, Iceland) or parallel coordinates plots to compare the two required key figures.

For each task we analytically determined a reference solution by counting the number of tablets that are necessary to display all relevant data points simultaneously for answering the given task without extensive memory load. Instead, the visualizations would serve as external memory to reduce demand on human memory [20]. These reference solutions for tasks 1-3 had an optimal number of tablets of $N_{T1} = 4$, $N_{T2} = 5$, and $N_{T3} = 6$. Furthermore, the tasks were designed with an increasing task complexity, which was approximated according to the guideline of Wood [55]. We created two task sets with equivalent task complexities for within-subjects counter-balancing. Each pair had to perform three tasks per condition consisting of 10 questions in total. For each question they had to agree on an answer.

Apparatus

Figure 1A shows the setup of our study. It consisted of an office table (0.9 m wide) and six Apple iPad Air 2 (9.7''). We attached felt pads to the bottom of each tablet for smoother sliding on the table to foster micro-mobility. The participants sat opposite of each other on office chairs simulating typical face-to-face collaboration in meetings. We opted against a side-by-side seating configuration because we observed during a pre-study that this resulted in much less simultaneous use of multiple tablets since groups tended to use only one shared tablet. The tablets were provided as a shared resource for both participants on a stack in the center of the table that was equally accessible from both sides. The participants were also provided with pen and paper for note taking.

Each session was recorded with a camera and microphone, so that participants' interactions and verbal communication became available for analysis. One experimenter was present in the study room and two further experimenters watched a live stream from the camera in an adjacent room.

Participants

To recruit the participants, we sent out invitations to the study through university mailing lists and advertised the study on social media. We recruited 24 participants (6 female, 18 male) with a mean age of 27.9 years ($SD = 7.5$, aged 19-49). The mean professional experience was 5.4 years ($SD = 6.6$, ranged between 0-25). The highest completed level of education for 8 participants was a high-school degree, for 7 a bachelor's degree, for 5 a Master's degree, and 3 held a PhD degree. One participant did not answer this question. 6 had a non-technical background. The 24 participants were randomly assigned and formed dyads resulting in 12 groups (6 groups were both male, 6 were mixed). All groups were systematically counterbalanced using Latin Square in relation to the two conditions and question set.

Procedure

The experimenter introduced participants to the study and received informed consent. Participants were asked to fill out a questionnaire regarding demographics and previous experience. Then they received an introduction to the system during a training phase with three example questions. During this, they were allowed to ask questions about the use of the system or task. Before the PAIRED condition they were also introduced to the advantages of device linking, and possible arrangements to work in mixed-focus collaboration [51]. To encourage efficient collaboration we asked them to finish each condition in 20 minutes. Nevertheless, we let every group finish all tasks. At the end of each condition, they were asked to individually fill out the NASA TLX questionnaire [17]. Lastly, we conducted a short semi-structured interview about their utilization of multiple devices and their preference for PAIRED vs. SINGLE and the number of tablets.

DATA ANALYSIS

We employed several data sources for data triangulation: (i) We recorded audio-visual data to analyze participants' verbal communication and their utilization of multiple tablets. (ii) We logged user interface selections, filter settings, and

tablet movements to quantify and analyze users' interactions with tablets. (iii) We asked users to fill out NASA TLX questionnaires to measure their subjective workload for each condition. (iv) We gathered demographic data and information about participants' previous experience with tablet use. (v) After each session the primary experimenter conducted a semi-structured interview to better understand the underlying barriers and biases of multi-tablet use. The two experimenters in the adjacent room, who had closely observed the sessions and collected a list of reoccurring themes during the study, also proposed additional questions for this interview via instant messaging. These concluding interviews enabled us to validate our initial interpretations of the observations. We then derived hypotheses for higher level findings and confirmed/rejected them by thematic coding of all video materials and, where possible, quantitative statistical analysis. For statistical data analysis, significance tests assume a p -value $< .05$ and non-parametric tests were used when the assumption of normal distribution was violated.

Video Data

To qualitatively and quantitatively analyze our video data, we developed a coding scheme. Thereby we focused mainly on participants' utilization of multiple tablets and the change between the number of used tablets. Additionally we also looked at other interesting and noteworthy behavior (e.g., deictic gestures, micro-movements). After initial viewings of our videos, we started to describe different codes that offered a sound description of the video data. This coding scheme was further refined until we felt confident that it met the requirements of our research goals. All video files were coded by one person to ensure consistency.

Coding Scheme

1.0 Tablet Utilization: To analyze how long and how much tablets each participant and group was using, we distinguish between the codes *1.1 Participant N is using Tablet P*, *1.2 Participant N is not using Tablet P*, *1.3 Both participants use Tablet P* and *1.4 Both participants are not using Tablet P*.

2.0 Task & Condition: To code the beginning of each condition and of each task, we used *2.1 Task X begins*, *2.2 Task X ends*, *2.3 Condition Y begins* and *2.4 Condition Y ends*.

3.0 Movements: For recording all noteworthy movements of tablets, we coded *3.1 Micro-Movement* and *3.2 Macro-Movement*. They were coded in addition to automated logging of tablet movements to also associate them with a particular participant.

4.0 Miscellaneous: To code all other interesting and conspicuous situations, we coded *4.1 Multi-tablet interaction*, *4.2 Reading Guide & Ruler* and *4.3 Pointing & Deictic references*.

Data Logs

In addition to the video recordings, we collected logs of user interactions with the tablets. We logged changes of visualization parameters (e.g., visualization type, country, year) and accelerometer movements of the tablets together with a timestamp and an internal ID of the tablet.

Time-weighted tablet utilization score (TUS)

To quantify the overall tablet utilization of each group as well as individual participant, we developed a time-weighted tablet utilization score (TUS). The mean TUS ranges between 1.0 (one tablet used for 100% of session) to 6.0 (six tablets used for 100% of session) and is calculated as follows: n is the number of utilizations, u is the number of currently used tablets, $t_{i+1} - t_i$ is the time for which u tablets are used, and $t_n - t_0$ is the total time of the task or condition:

$$TUS = \sum_{i=0}^{n-1} \frac{u_i \cdot (t_{i+1} - t_i)}{t_n - t_0} \quad (1)$$

We opted against an automatic determination of u based on log files, eye tracking, or similar and consider human judgement and interpretation of video data as indispensable. Typical visualization tasks with multiple devices include many shifts of visual attention for comparison or also staring off into space while thinking. Relying on automated analysis of eye tracking or logs alone would therefore lead to many misinterpretations of gaze or interactions.

FINDINGS & DISCUSSION

In the following, we report on our four findings about typical usage patterns, user strategies, and noteworthy exceptions. We also discuss their implications for the design of future multi-tablet visualization tools. For better readability, we highlight resulting numbers from statistical analysis in blue.

Finding 1: Inefficient Tablet Utilization & Legacy Bias

All groups successfully completed their tasks. The mean total completion time for all tasks in both conditions was 37.8 min per group ($SD = 8.3$ min, $Min = 24.6$ min, $Max = 52.4$ min). On average, the groups correctly answered 18.75 ($SD = 1.14$) of 20 questions. This confirms that the tasks were solvable and neither too simple or difficult.

Inefficient tablet utilization

A key observation was that participants did not simultaneously use multiple tablets as often as we expected and much less than the we expected the tasks would afford. This was the case even though the participants were instructed in using two tablets or more during the study. Although our tasks were designed to clearly favor the simultaneous use of two tablets or more *per user*, the actual mean TUS during all tasks and conditions was only 2.06 ($SD = 0.23$) *per group*. For all tasks the minimum TUS was 1.1 and the maximum TUS 3.5. Despite a few brief incidents during which 4 of 6 tablets were used, the overall simultaneous use of tablets remained much smaller than what would have been theoretically possible. It was also much smaller than in our reference solutions which would have been less demanding in terms of necessary user interactions and users' memory load. These reference solutions for tasks 1-3 involved between 4 and 6 tablets per group ($N_{T1} = 4$, $N_{T2} = 5$, $N_{T3} = 6$). In contrast, the actual mean TUS per group for tasks 1-3 during the study were: $\overline{TUS}_{T1} = 1.96$ ($SD = 0.39$, $Min = 1.45$, $Max = 2.60$), $\overline{TUS}_{T2} = 2.09$ ($SD = 0.23$, $Min = 1.80$, $Max = 2.50$), $\overline{TUS}_{T3} = 2.14$ ($SD = 0.34$, $Min = 1.75$, $Max = 2.80$).

Legacy bias: "Computers" vs. "Documents"

We attribute this inefficient tablet utilization to the phenomenon of "legacy bias" that was first described by Morris et al. for gesture elicitation studies [35]. It means that users are often biased by their experience with prior interfaces and technologies, particularly the WIMP (windows, icons, menus, and pointing) interfaces that have been standard on traditional PCs for the past two decades [35].

We observed a similar bias in our study for how participants used tablets for collaborative visualization. Participants hesitated to use tablets as "documents", i.e., as physical manifestations or snapshots of selected content such as important data views or intermediate results. Instead, participants typically used only very few tablets in parallel and as if they were multi-purpose "computers" similar to traditional desktop PCs. Tablets were used to access data *through* them rather than to contain data *in* them. Participants also constantly changed the function and content of the "computers" by interacting with them and did not assign them with the longer lasting role or content of a "document". This legacy bias towards using tablets like "computers" rather than "documents" was also confirmed by the participants themselves: A member of G7 mentioned that it felt natural to only use as many tablets during the study as he has screens connected to his desktop PC. A member of G4 commented that his usual way of utilizing the screens of his desktop PC served him as a point of reference for using multiple tablets during the study. The other member of G7 mentioned that he usually works only with one screen and that he therefore felt that using more than one tablet would not have made any sense during the study. Furthermore, members of the groups G3/G8/G9 said that using more than one or two tablets is unfamiliar and overburdening.

In contrast, if participants had followed the original Weiserian vision of tablets/pads as paper-like commodities, they should have used many more in parallel. Similar to how we organize papers on our desks [32], each tablet could have served as a single-purpose "document" at a meaningful position, containing only selected data in a particular visual representation. Moreover, a tablet could have been put aside at anytime to keep it as physical "snapshot" or "bookmark" of intermediate results to then continue working with another tablet. Such a distribution of content and roles across many tablets as "documents" would also reduce the amount of interactions that are needed to navigate back and forth between different views when tablets are used as "computers".

Our distinction between "computers" and "documents" is intended to illustrate two opposing roles that tablets can play in multi-tablet computing. However, they should not be seen as mutually exclusive. Depending on the users' intentions and the task at hand, tablets can switch between these roles or even combine aspects of both at the same time. We believe that introducing this duality of roles can help researchers and designers to better analyze and reason about how users will perceive and use future multi-tablet system.

Reducing legacy bias

Previous work on gesture elicitation studies, showed that legacy bias can be reduced by letting users work in groups

and intensifying the *priming* phase, i.e. demonstrating innovative ways of using the target technology to users in order to inspire them and let them think more generally about how to accomplish a given task [35, 40]. In our case, neither working in groups of two nor demonstrating the benefit of multiple tablets (e.g. the experimenter demonstrated the benefit of linked tablets with multiple coordinated views before the PAIRED condition) seemed to have a significant effect on tablet utilization. Since 23 of 24 participants reported that they are familiar with tablets and how to use them, this cannot be explained by a simple lack of experience with using tablets. Furthermore, our design of the tasks clearly favored using multiple tablets, so that over time a gradual learning process could have taken place. Nonetheless the extent of tablet utilization remained small and below that of our reference solutions for all tasks. There was also no significant difference in mean TUS per group between SINGLE and PAIRED, $t(11) = -1.911$, $p = .082$, $r = .50$.

Comparison to related studies

Our results are noteworthy with regard to the utilization study of the “Conductor” system by Hamilton et al. [16]. Although the sensemaking task in their study was remotely similar to ours, they observed a typical usage of 5 to 10 tablets for their single-user multi-tablet sensemaking tool. We believe that their very different findings should not be considered a contradiction to our results and that they are explicable by Conductor’s specific system design and study tasks. Conductor is specifically designed for *easy functionality chaining* [16] rather than using parallel views. Multiple devices are necessary to create a functionality chain of sufficient complexity to complete tasks. Nonetheless Hamilton et al. also mention how one user considered the tablets as physical equivalents to an application window in a desktop UI [16]. This seems to be similar to what we call “legacy bias”.

Another related study is by Chen et al. who provided graduate students in the humanities with four linked e-reader tablets and observed how they supported active reading tasks during month-long deployments in the wild [10]. They found that no users carried around more than two tablets and that two tablets were sufficiently helpful for many tasks such as viewing two parts of the same document simultaneously without the need for traditional window management. Chen et al. also mention that ironically the advantage of a multi-tablet system is that it also gave participants the freedom to use only a single slate. For many activities this was sufficient and often more desirable since it eliminated superfluous hardware and increased portability [10]. Even though there are great differences between Chen et al. and our application domain and system design, their observations resonate with our experiences during our study and that using two views on the same data was most popular and appeared sufficient to the participants most of the time.

Design implications

We cannot predict how the tablet utilization would have developed during a longer observation than in our study. But there are some indicators such as the absence of a significant increase in the TUS between the first condition and the second

condition ($Mdn_{First} = 1.95$ and $Mdn_{Second} = 2.03$, $z = -.356$, $p = .748$, $r = -0.10$), different comments from users, and the fact that Chen et al. did not observe more utilization during a month-long deployment of a multi-tablet system. We therefore believe that the observed underutilization of multiple tablets is not just an artifact of our tasks, study, system design, or sample of the user population. Instead, we believe it is more deeply rooted in users’ prior experiences of using WIMP interfaces. This constitute a legacy bias towards using tablets as “computers” rather than “documents”.

As designers of future multi-tablet systems, we should therefore think about ways to increase the affordance of multi-tablet systems for multi-tablet use to overcome this bias. Even when their benefit is quite obvious (like in our study), cross-device use and parallel views are not self-explanatory. Therefore designers should consider actively exposing users to more and more advanced functionality and explaining their uses and benefits. Similar approaches have been proposed as *multi-layered user interfaces* [50] in the context of information visualization for gradually teaching users about the benefits of multiple coordinated views. Other strategies include to support *intermodal improvement* and *vocabulary extensions* [11], e.g., by improving the learnability of mobile devices [26] or executable context-specific action suggestions [13].

Finding 2: Benefits of and Barriers for Multi-Tablet Use

To better understand the observed legacy bias and what kept participants from using multiple tablets, we first identified the most frequent tablet configurations. To specify configurations we use a simple notation with the two numbers of tablets used by participant A and B in parentheses separated by a colon, e.g., “(A:B)”. The three most frequent utilization patterns over all groups, tasks, and conditions were “one tablet per user (1:1)” (66.8% of the time), “one tablet per group (1:0 | 0:1)” (12.9% of the time), and “one tablet and two tablets (1:2 | 2:1)” (10.2% of the time). Please note that for our configurations we only counted the tablets that were actively used or observed by users and not tablets that were placed somewhere on the desk but not actively used or observed.

	One tablet per user. (1:1)	One tablet per group. (1:0 0:1)	One tablet and two tablets. (1:2 2:1)
SINGLE	73.7%	17.1%	4.7%
PAIRED	59.3%	11.6%	15.0%
BOTH	66.8%	12.9%	10.2%

Table 1. The three most frequent tablet configurations per condition.

Additionally, we analyzed the number of transitions between different configurations. Figure 2 shows a transition graph with the number of transitions between selected configurations. Please note that for space and readability reasons we do not show a complete graph with all configurations but only with the most frequent transitions.

One tablet per user (1:1)

This was the most frequent configuration regardless of whether tablets were linked (PAIRED) or not (SINGLE).




Figure 2. Transition graph showing most frequent transitions between configurations. The number on each transition indicate the counted number of transitions between two configurations.

It was also the most frequent starting configuration for all groups and tasks (36 of 72 tasks: (50.0%). Since both participants were asked to complete the task as a team, they often began with one tablet for each person regardless of the specific task. Participants confirmed this by commenting that that one tablet per user felt “natural” and “familiar”. Only after the nature of the task was better understood through personal experimentation on each tablet, participants started to systematically involve more tablets/views. For all groups in both conditions there were $N_{more} = 48$ transitions from (1:1) to “more tablets” and $N_{less} = 36$ transitions from “more tablets” to (1:1). This results in a return rate of $N_{less}/N_{more} = 75\%$ and hints at the role of this configuration as a starting point from which more complex uses with more tablets evolved. However, as becomes visible in the TUS scores from Finding 1 and the return rate, the great majority of these uses were only temporary. Users seemed to be willing to try to use multiple tablets, but did not succeed in using them efficiently. This also showed when G5 and G12 commented that they would now use all tablets for their tasks and placed them accordingly (e.g. Figure 1A) but even then they quickly returned to using as few tablets as before.

Furthermore, this configuration almost exclusively served *loosely coupled* [51] parallel work with two unlinked views instead of *tightly coupled* [51] collaboration across linked views. In 72 of 87 cases (79.2%), the tablets did not belong to the same pair so that there was no functional connection at all. In the other 15 of 87 cases (20.8%), it was only G5/G6/G7/G12 who picked tablets from the same pair with connected views but without exposing a clearly identifiable goal or strategy. It seemed that tablets were more or less randomly picked from the stack of available tablets and that participants hardly noticed if tablets were or were not linked. If they were, participants started to verbally coordinate their steps to avoid that the interactions by the other team member disturbed own work or vice versa, but there was no clear evidence for a deliberate and more coordinated use of cross-device functionality in the (1:1) configuration. Similarly, we only observed few cases, in which participants closely worked together by putting their devices side-by-side, e.g., to visually compare them with or without a functional connection. This only happened in 15 of 87 cases (17.2%).




Figure 3. “One tablet per group” with tablet oriented towards the participant on the right (A) and rotated by 90 degrees to facilitate reading from both sides (B).

One tablet per group (0:1 | 1:0)

This configuration with only one tablet per group was the second most frequent configuration (see Table 1). It was also the second most frequent starting configuration for all groups and tasks (32 of 72 tasks: 44.4%). It enabled users to engage in tightly-coupled face-to-face collaboration and provided a shared surface for deictic and verbal references similar to an interactive tabletop (see for example [51]). Figure 3A shows an example situation where both participants focused on a single tablet in the center although four tablets were idling in the stack and a fifth tablet was right in front of one participant. It is noteworthy that this tablet was already linked with the center tablet, so that the views on both tablets were synchronized. Nonetheless, participants preferred to use only one of the two linked tablets, although that meant that one of them had to read upside down.

In general, we observed that participants often struggled with reading orientation and tried to adjust the orientations of shared tablets to accommodate both group members. Since this is impossible many participants rotated tablets by 90 degrees as a trade-off. This strategy of avoiding exclusion can be considered as thoughtful and polite, but it also resulted in leaning on the table or visibly uncomfortable seating and head positions for both members (Figure 3B).

Since many groups also began their tasks with this configuration, transitions between this configuration to the popular (1:1) were most frequent. Their frequency can be explained by the popularity of both configurations and the need for transitions between tightly coupled and loosely coupled collaboration during typical *mixed-focus collaboration* [51]. As the archetypes for tightly and loosely coupled work, both configurations served as starting points and attractors during collaboration and therefore their role and that of transitions between them was central to group work.

Multiple tablets per user (A:B with A|B > 1)

The use of more than one tablet/view per user (example in Figure 1A) typically served the comparison or relation of values for different years or countries across different tablets or visualization types. For example, one member of G2 mentioned that the parallel utilization of multiple tablets was very useful for comparisons of two data points or to provide different views of the same data point. This resonates with the original idea of multiple coordinated views to let users understand their data better by enabling them to interact with the

presented information and view it through different representations [41]. The tablets to compare were typically arranged side-by-side with two or more tablets in a row or column to reduce the need for head or eye movement when switching attention between different visual items.

Another interesting observation was that some participants used multiple coordinated views in the PAIRED condition to adjust the settings of their first tablet using a second tablet. A member of G12 explained his reasons for this by mentioning that he considered the tablets as a kind of tangible screens and used one tablet as “configurator” for the different settings (e.g. visualization, filters) of other tablets.

SINGLE vs. PAIRED tablets

Unlike we initially expected, there were no fundamental differences between the collaboration styles or the use of tablets in the SINGLE and PAIRED condition. Although 19 of 24 participants answered that they preferred PAIRED, there was also no statistically significant effect of SINGLE vs. PAIRED on the task completion time ($Mdn_{SINGLE} = 19.0$, $Mdn_{PAIRED} = 20.4$, $z = -1.334$, $p = .204$, $r = -0.39$). However, some differences for the configuration are visible in Table 1: There was a significantly lower usage of (1:1) in the PAIRED condition ($Mdn_{1:1} = .631$) than in the SINGLE condition ($Mdn_{1:1} = .739$), $z_{1:1} = -1.96$, $p < .05$, $r = -.57$. Also 2:1 and 2:2 were significantly higher in the PAIRED condition ($Mdn_{2:1} = .146$, $Mdn_{2:2} = .006$) than in the SINGLE condition ($Mdn_{2:1} = .140$, $Mdn_{2:2} = .000$), $z_{2:1} = -2.26$, $p < .05$, $r = -.65$ and $z_{2:2} = -2.03$, $p < .05$, $r = -.59$.

These results can be interpreted as evidence for the participants’ subjective preference and interest in using multiple views and cross-device functionality that showed in an increased use of more than one tablet per person in the PAIRED condition but did not pay off in terms of significantly better results or task completion times. The overall positive attitude towards paired tablets became also evident in comments during the concluding interview. G6/G8/G12 commented that this functionality motivated to work more closely together. G12 also mentioned that it helped to be able to transfer settings between devices. However, participants were also critical and found the functionality in PAIRED demanding. G8 commented that working with paired tablets felt very odd and G10/G11 suggested that pairing always should be optional. G2/G5/G7 found that they only realized the potential benefit after some time and being more familiar with the system.

This more critical view also resonates with the analysis of the self-reported mental demand of using multiple tablets. A Wilcoxon signed-rank test revealed statistical significant differences for the NASA TLX in the overall subjective workload (overall measure of all subscales) and individually in the mental demand subscale. The subjective workload was significantly higher in the PAIRED condition ($Mdn = 43.33$) than in the SINGLE condition ($Mdn = 37.08$), $z = -3.615$, $p < .001$, $r = -.74$. Similarly, users also reported a significantly higher mental demand in the PAIRED condition ($Mdn = 70.00$) than in the SINGLE condition ($Mdn = 65.00$), $z = -2.251$, $p < .05$, $r = -.50$.

In conclusion, while multiple views and cross-device interactions attract users’ attention and interest, they can introduce a higher mental demand and do not necessarily improve efficiency and effectiveness. While users see great potential for using multiple views across devices, our design and implementation did not achieve measurable benefits.

Design implications

The typical starting configurations were “one tablet per user” for loosely coupled parallel work and “one tablet per group” for tightly coupled collaboration. They served as natural entry points from which the users started their visual exploration of the data space using more (or sometimes also less) tablets. Therefore these entry points should be designed with care and should support users in quickly discovering the potential benefits of more complex configurations and also how to switch between different coupling styles. Furthermore, participants should be immediately rewarded when using more tablets instead of being discouraged by a higher mental demand.

Finding 3: Use of Space

Space is an invaluable resource to exploit in order to facilitate everyday problem solving and planning [25]. It can serve us as externalization to simplify choice, perception, and internal computation [25]. Space can be used to encode ordering information and to off-load memory by creating a persistent physical reference [19]. HCI has therefore extensively discussed the use of desks, rooms, or large screens to spatially organize digital items or mobile devices as physical containers of digital content, e.g., [2, 16, 25, 32, 39, 49]. For example, prior studies reported about “incremental formalizations” [49] during sensemaking, i.e., how users evolved configurations of items in space over time from simple stacks or clusters after rough categorizations towards more formal and semantically meaningful configurations, e.g., a timeline of objects that is ordered by date [2]. Since the role of space can be manifold, we here distinguish its potential benefits between using tablets as “computers” vs. “documents”.

Space for “computers”: Ergonomics and collaboration

When using tablets as “computers”, users can benefit from space by quickly creating more ergonomic configurations of devices that better fit into the physical or social setting, e.g., increasing readability and reachability of the tablets, supporting face-to-face collaboration. For example, in Chen et al.’s study, tablets provided a more tailored use of space, which offered flexibility over a variety of settings and for many different tasks [10]. Chen et al. observed users putting tablets on their lap while working with a laptop or they put one on top of the other while reading documents as passengers in a car during a road trip.

We observed similar benefits of space for “physical window management”. A user’s central workspace (or personal territory [48]) could be quickly reconfigured according to the task at hand, simply by rearranging tablets without the need to grab or drag virtual handles, title bars, or using snapping mechanisms. For example, it was easy to create more space for an additional tablet with simple physical movements (Figure 4). As illustrated in Figure 5B, it was also possible to (at least temporarily) use the space above tables



Figure 4. Physical reconfiguration of workspace.



Figure 5. Tilting tablets to increase readability, facilitate comparisons, and use space above the table.

or desks by tilting devices where there would not be enough space otherwise. By this, users were also able to easily put or hold tablets side-by-side to facilitate visual comparisons (Figure 4B,5B,6B). Such spatial configurations to facilitate comparisons between two or more views were used 10 times during the study. In 6 cases the tablets were placed in a row, meaning side-by-side to the left and right edges of their screens. In 4 cases the tablets were placed in a column, meaning side-by-side to the top and bottom edges of their screens.

Using space for improving ergonomics did not only help single users but also the teams. We frequently observed how users moved tablets from their personal territory to group territories [48] by rotating or tilting tablets towards their collaborators in order to improve readability or highlight content to share (Figure 5A). In contrast to flat interactive surfaces such as large vertical screens or interactive tabletops, the observed micro-mobility [31, 33] of tablets enabled very fast and flexible changes of spatial configurations that accommodated different collaboration styles, even if screen orientation remained to be an issue throughout the study (see Figure 3B).

Space for “documents”: Piling and filing

When using tablets as “documents”, users can benefit from space by creating meaningful clusters or semantic spatial configurations of tablets. Different locations can be used to store or categorize content. For example, Hamilton et al. [16] report that all participants of their study made extensive use of physical organization of tablets to keep track of information. Chen et al. [10] report that some users used tablets to open documents and put them aside as physical reminders of things they still needed to read.

We also witnessed comparable uses, e.g., putting tablets in certain configurations or orders as reminders of what to compare or as externalizations of thought processes. In one case,



Figure 6. Reordering two tablets by year.

G1 realised that they had not completely answered the previous question before starting to work on the next question. Therefore they interrupted their work and put aside the currently used tablet as a reminder (similar as reported in [10]) and container for intermediate results to then switch to work on the previous question with another tablet.

We also observed behavior similar to the afore-mentioned “incremental formalizations” [2, 49]. For example, one member of G5 re-arranged his two tablets that contained views of data from different years to order them by year from left to right according to his familiar reading direction (Figure 6). We observed similar behavior of establishing semantically meaningful orders in 4 cases.

However, in comparison to the use of virtual space in the study by Andrews et al. [2] or the typical organization of printed documents on office desks [32], we did not observe a more permanent placement of tablets in piles, clusters, or stacks based on their content. Except the stack in which the tablets were provided to participants at the beginning of each condition, there was no physical piling, filing, or clustering of tablets. Since the use of tablets as “computers” dominated during the study, a more permanent spatial organization of tablets as “documents” did not take place.

Design implications

The observed spatial configurations of the tablets carried meaning and were not results of random movements. For example, users actively reordered multiple views in space according to familiar conventions, e.g., ordering them by year from left to right. They also put tablets side-by-side for comparisons or handed or rotated them to other users for collaboration. In future, making tablets *spatially-aware* [39, 56] (i.e., aware of their positions and mutual spatial relations) can be used to proactively react to such movements. For example, users can be provided with spatially-aware menus and modes [39] and automatic adjustments of visualizations, settings, couplings, or content according to the assumed user goal.

Finding 4: Physical tools and techniques

Another advantage of using multiple tablets was the possibility to physically point at, easily grab, or touch a tablet’s display. This enabled users to transfer familiar physical techniques and tools from working with paper documents to working with the digital “documents” on a tablet. These tools and techniques supported collaboration as well as reading and comprehending visualizations.




Figure 7. Using a sheet of paper or hand as visual guide.

For example, participants used pens or their hands to perform deictic gestures during the sensemaking process and for discussion during collaboration. Deictic gesturing was mostly used to draw the attention of the other team member to a specific data point, visualization, or tablet. It was also used to ask the other participant to change the role of a tablet, e.g., pointing at a tablet with the wish to use it or change the visualization. Similar to interactive tabletops, our multi-tablet system for co-located collaboration created a shared space and frame of reference in which natural deictic gestures and face-to-face collaboration became possible.

Furthermore, also less obvious physical techniques and tools were applied by users that would not have been possible when using conventional vertical screens (e.g., laptop, PC screen). Users frequently used sheets of paper, pens, or their hands as reading guides, improvised rulers, or to cover irrelevant information. As seen in Figure 7A, a piece of paper was used as a ruler to visually compare the vertical position of data points in the parallel coordinates view. Figure 7B shows an example where a participant used her hand to mentally divide a map into two regions to visually focus on the relevant one. Another frequently used technique was entirely covering irrelevant parts of the screen with paper for reducing the amount of visual information to process and by this simplifying perception and internal computation.

Design implications

This finding can inform the design of future multi-tablet systems in different ways. First of all, it shows what digital functions or visual elements were missing (e.g. visual guides or rulers) and should be considered in future. On the other hand, it showed that users become inventive and are also able to compensate for missing functionality as soon as they have sufficient physical access to the screen. Therefore not all functionality must be provided digitally. However, designers should make their designs resilient against unintended touch events that might result from physically accessing the tablets and using hands or objects on the screen.

LIMITATIONS AND FUTURE WORK

The design of the system and task naturally influenced the behavior that we observed. Firstly, the tasks were designed as visualization and sensemaking tasks, limiting the generalizability of our results to arbitrary computing activities. Secondly, some subtasks made it necessary to collaboratively interpret and discuss geographic correlations and therefore strongly shifted the focus to verbal communication between the participants. Thirdly, the participants had only a short

time to get familiar with the visualizations, tasks, and the cross-device functionality. Despite the fact that we gave every group a comprehensive introduction to the system and its functions, some participants might have needed more time to get used to the system. However, as we discuss in finding 1, there were no indications for a significant learning effect between the first and second condition during our study. Finally, the group size could have influenced the utilization of tablets per user. A group size of three or more participants might have afforded more simultaneous use of tablets.

Our study shows that even when our participants were provided with tablets in abundance, and simple means for exploiting cross-device functionality, the tablets were not used as we had anticipated. We have attributed this to a “legacy bias” that originates from cultural upbringing with personal computers. Yet, our study also indicates that cross-device functionality has a higher cognitive demand than singular use of devices. A future study with a similar task but after participants had extensive training in using multiple tablets and cross-device functionality could reveal if the TUS would differ and the mental demand would be lower. Furthermore, if cross-device functionality should become a common feature of consumer technology, we could verify the “legacy bias” hypothesis. Finally, we have exclusively studied a sensemaking and visualization task. In future, a similar hardware setup could be used to also study other tasks, e.g., content production, remote collaboration, or monitoring external systems.

CONCLUSION

The Weiserian vision of users conducting knowledge work seamlessly across multiple mobile devices has been a driving force in HCI research on cross-device interaction. Today, however, our tablets play a quite different role than the pads Weiser envisioned. As a consequence, using many tablets at a time and sharing them with collaborators rarely happens outside research labs. We therefore conducted a controlled experiment to contribute to a better understanding of the benefits and barriers of multi-tablet use during collaborative information visualization tasks.

Findings of our study show that participants underutilized the available tablets and were hesitant to use multiple tablets in parallel, hinting at a “legacy bias” of using tablets as “computers” rather than “documents”. Also, while multiple coordinated views and cross-device interactions attracted users attention and interest, they introduced a higher mental demand and did not improve efficiency and effectiveness. Furthermore, movements of tablets in space served (i) to improve ergonomics and collaboration and (ii) to establish semantically meaningful orders. However, we did not observe physical piling, filing, or clustering of tablets based on content. Finally, participants transferred natural techniques and tools from working with paper to working with the digital “documents” on a tablet, e.g., using hands or paper as guides or rulers, revealing new design opportunities for multi-tablet visualization systems.

REFERENCES

1. Gregory D. Abowd. 2012. What Next, Ubicomp?: Celebrating an Intellectual Disappearing Act. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing (UbiComp '12)*. ACM, New York, NY, USA, 31–40. DOI : <http://dx.doi.org/10.1145/2370216.2370222>
2. Christopher Andrews, Alex Endert, and Chris North. 2010. Space to Think: Large High-resolution Displays for Sensemaking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, USA, 55–64. DOI : <http://dx.doi.org/10.1145/1753326.1753336>
3. Matthew P. Aylett and Shaun Lawson. 2016. The Smartphone: A Lacanian Stain, A Tech Killer, and an Embodiment of Radical Individualism. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems - CHI EA '16*. ACM Press, New York, New York, USA, 501–511. DOI : <http://dx.doi.org/10.1145/2851581.2892581>
4. Matthew P. Aylett and Aaron J. Quigley. 2015. The Broken Dream of Pervasive Sentient Ambient Calm Invisible Ubiquitous Computing. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems - CHI EA '15*. ACM, New York, NY, USA, 425–435. DOI : <http://dx.doi.org/10.1145/2702613.2732508>
5. Genevieve Bell and Paul Dourish. 2006. Yesterday's tomorrows: notes on ubiquitous computing's dominant vision. *Personal and Ubiquitous Computing* 11, 2 (nov 2006), 133–143. DOI : <http://dx.doi.org/10.1007/s00779-006-0071-x>
6. R. Borovoy and B. Knep. 2011. Junkyard Jumbotron, <http://jumbotron.media.mit.edu>. (2011).
7. Michael Bostock, Vadim Ogievetsky, and Jeffrey Heer. 2011. D3: Data-Driven Documents. *IEEE Trans. Visualization & Comp. Graphics (Proc. InfoVis)* (2011). <http://vis.stanford.edu/papers/d3>
8. Thorsten Buering, Jens Gerken, and Harald Reiterer. 2006. User Interaction with Scatterplots on Small Screens - A Comparative Evaluation of Geometric-Semantic Zoom and Fisheye Distortion. *IEEE Transactions on Visualization and Computer Graphics* 12, 5 (Sept. 2006), 829–836. DOI : <http://dx.doi.org/10.1109/TVCG.2006.187>
9. Nicholas Chen, Francois Guimbretiere, and Abigail Sellen. 2012. Designing a Multi-slate Reading Environment to Support Active Reading Activities. *ACM Trans. Comput.-Hum. Interact.* 19, 3, Article 18 (Oct. 2012), 35 pages. DOI : <http://dx.doi.org/10.1145/2362364.2362366>
10. Nicholas Chen, François Guimbretière, and Abigail Sellen. 2013. Graduate student use of a multi-slate reading system. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13*. ACM Press, New York, New York, USA, 1799. DOI : <http://dx.doi.org/10.1145/2470654.2466237>
11. Andy Cockburn, Carl Gutwin, Joey Scarr, and Sylvain Malacria. 2014. Supporting Novice to Expert Transitions in User Interfaces. *ACM Comput. Surv.* 47, 2, Article 31 (Nov. 2014), 36 pages. DOI : <http://dx.doi.org/10.1145/2659796>
12. Steven M. Drucker, Danyel Fisher, Ramik Sadana, Jessica Herron, and m.c. schraefel. 2013. TouchViz: A Case Study Comparing Two Interfaces for Data Analytics on Tablets. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 2301–2310. DOI : <http://dx.doi.org/10.1145/2470654.2481318>
13. C. Ailie Fraser, Mira Dontcheva, Holger Winnemöller, Sheryl Ehrlich, and Scott Klemmer. 2016. DiscoverySpace: Suggesting Actions in Complex Software. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems (DIS '16)*. ACM, New York, NY, USA, 1221–1232. DOI : <http://dx.doi.org/10.1145/2901790.2901849>
14. Jens Grubert, Matthias Kranz, and Aaron Quigley. 2016. Challenges in mobile multi-device ecosystems. *mUX: The Journal of Mobile User Experience* 5, 1 (2016), 1–22. DOI : <http://dx.doi.org/10.1186/s13678-016-0007-y>
15. Jonathan Haber, Miguel a. Nacenta, and Sheelagh Carpendale. 2014. Paper vs. tablets: the effect of document media in co-located collaborative work. In *Proceedings of the 2014 International Working Conference on Advanced Visual Interfaces - AVI '14*. ACM Press, New York, New York, USA, 89–96. DOI : <http://dx.doi.org/10.1145/2598153.2598170>
16. Peter Hamilton and Daniel J. Wigdor. 2014. Conductor: Enabling and Understanding Cross-device Interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 2773–2782. DOI : <http://dx.doi.org/10.1145/2556288.2557170>
17. Sandra G. Hart. 2006. Nasa-Task Load Index (NASA-TLX); 20 Years Later. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 50, 9 (oct 2006), 904–908. DOI : <http://dx.doi.org/10.1177/154193120605000909>
18. Ken Hinckley. 2003. Synchronous Gestures for Multiple Persons and Computers. In *Proceedings of the 16th Annual ACM Symposium on User Interface Software and Technology (UIST '03)*. ACM, New York, NY, USA, 149–158. DOI : <http://dx.doi.org/10.1145/964696.964713>
19. James Hollan, Edwin Hutchins, and David Kirsh. 2000. Distributed Cognition: Toward a New Foundation for Human-computer Interaction Research. *ACM Trans. Comput.-Hum. Interact.* 7, 2 (June 2000), 174–196. DOI : <http://dx.doi.org/10.1145/353485.353487>

20. Weidong Huang, Peter Eades, and Seok-Hee Hong. 2009. Measuring Effectiveness of Graph Visualizations: A Cognitive Load Perspective. *Information Visualization* 8, 3 (June 2009), 139–152. DOI : <http://dx.doi.org/10.1057/ivs.2009.10>
21. Petra Isenberg, Danyel Fisher, Sharoda A. Paul, Meredith Ringel Morris, Kori Inkpen, and Mary Czerwinski. 2012. Co-Located Collaborative Visual Analytics around a Tabletop Display. *IEEE Trans. Vis. Comput. Graph.* 18, 5 (2012), 689–702. DOI : <http://dx.doi.org/10.1109/TVCG.2011.287>
22. Mikkel R. Jakobsen and Kasper Hornbæk. 2014. Up Close and Personal: Collaborative Work on a High-resolution Multitouch Wall Display. *ACM Trans. Comput.-Hum. Interact.* 21, 2, Article 11 (Feb. 2014), 34 pages. DOI : <http://dx.doi.org/10.1145/2576099>
23. Hans-Christian Jetter, Jens Gerken, Michael Zöllner, Harald Reiterer, and Natasa Milic-Frayling. 2011. Materializing the Query with Facet-streams: A Hybrid Surface for Collaborative Search on Tabletops. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 3013–3022. DOI : <http://dx.doi.org/10.1145/1978942.1979390>
24. Tero Jokela, Jarno Ojala, and Thomas Olsson. 2015. A Diary Study on Combining Multiple Information Devices in Everyday Activities and Tasks. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 3903–3912. DOI : <http://dx.doi.org/10.1145/2702123.2702211>
25. David Kirsh. 1995. The Intelligent Use of Space. *Artif. Intell.* 73, 1-2 (Feb. 1995), 31–68. DOI : [http://dx.doi.org/10.1016/0004-3702\(94\)00017-U](http://dx.doi.org/10.1016/0004-3702(94)00017-U)
26. Rock Leung, Leah Findlater, Joanna McGrenere, Peter Graf, and Justine Yang. 2010. Multi-Layered Interfaces to Improve Older Adults’ Initial Learnability of Mobile Applications. *ACM Trans. Access. Comput.* 3, 1, Article 1 (Sept. 2010), 30 pages. DOI : <http://dx.doi.org/10.1145/1838562.1838563>
27. Ming Li and Leif Kobbelt. 2012. Dynamic tiling display: building an interactive display surface using multiple mobile devices. In *Proceedings of the 11th International Conference on Mobile and Ubiquitous Multimedia - MUM '12*. ACM Press, New York, New York, USA, 1. DOI : <http://dx.doi.org/10.1145/2406367.2406397>
28. Andrés Lucero, Jussi Holopainen, and Tero Jokela. 2011. Pass-them-around: Collaborative Use of Mobile Phones for Photo Sharing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 1787–1796. DOI : <http://dx.doi.org/10.1145/1978942.1979201>
29. Andrés Lucero, Matt Jones, Tero Jokela, and Simon Robinson. 2013. Mobile Collocated Interactions: Taking an Offline Break Together. *interactions* 20, 2 (March 2013), 26–32. DOI : <http://dx.doi.org/10.1145/2427076.2427083>
30. Andrés Lucero, Jaakko Keränen, and Hannu Korhonen. 2010. Collaborative Use of Mobile Phones for Brainstorming. In *Proceedings of the 12th International Conference on Human Computer Interaction with Mobile Devices and Services (MobileHCI '10)*. ACM, New York, NY, USA, 337–340. DOI : <http://dx.doi.org/10.1145/1851600.1851659>
31. Paul Luff and Christian Heath. 1998. Mobility in Collaboration. In *Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work (CSCW '98)*. ACM, New York, NY, USA, 305–314. DOI : <http://dx.doi.org/10.1145/289444.289505>
32. Thomas W. Malone. 1983. How Do People Organize Their Desks?: Implications for the Design of Office Information Systems. *ACM Trans. Inf. Syst.* 1, 1 (Jan. 1983), 99–112. DOI : <http://dx.doi.org/10.1145/357423.357430>
33. Nicolai Marquardt, Ken Hinckley, and Saul Greenberg. 2012. Cross-device Interaction via Micro-mobility and F-formations. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology (UIST '12)*. ACM, New York, NY, USA, 13–22. DOI : <http://dx.doi.org/10.1145/2380116.2380121>
34. Will McGrath, Brian Bowman, David McCallum, Juan David Hincapié-Ramos, Niklas Elmquist, and Pourang Irani. 2012. Branch-explore-merge: Facilitating Real-time Revision Control in Collaborative Visual Exploration. In *Proc. of the 2012 ACM International Conference on Interactive Tabletops and Surfaces (ITS '12)*. ACM, New York, NY, USA, 235–244. DOI : <http://dx.doi.org/10.1145/2396636.2396673>
35. Meredith Ringel Morris, Andreea Danilescu, Steven Drucker, Danyel Fisher, Bongshin Lee, m. c. schraefel, and Jacob O. Wobbrock. 2014. Reducing Legacy Bias in Gesture Elicitation Studies. *Interactions* 21, 3 (May 2014), 40–45. DOI : <http://dx.doi.org/10.1145/2591689>
36. Meredith Ringel Morris, Danyel Fisher, and Daniel Wigdor. 2010a. Search on Surfaces: Exploring the Potential of Interactive Tabletops for Collaborative Search Tasks. *Inf. Process. Manage.* 46, 6 (Nov. 2010), 703–717. DOI : <http://dx.doi.org/10.1016/j.ipm.2009.10.004>
37. Meredith Ringel Morris, Jarrod Lombardo, and Daniel Wigdor. 2010b. WeSearch: Supporting Collaborative Search and Sensemaking on a Tabletop Display. In *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work (CSCW '10)*. ACM, New York, NY, USA, 401–410. DOI : <http://dx.doi.org/10.1145/1718918.1718987>
38. Antti Oulasvirta. 2008. When users ”do” the Ubicomp. *interactions* 15, 2 (mar 2008), 6. DOI : <http://dx.doi.org/10.1145/1340961.1340963>

39. Roman Rädle, Hans-Christian Jetter, Nicolai Marquardt, Harald Reiterer, and Yvonne Rogers. 2014. HuddleLamp: Spatially-Aware Mobile Displays for Ad-hoc Around-the-Table Collaboration. In *Proceedings of the Ninth ACM International Conference on Interactive Tabletops and Surfaces (ITS '14)*. ACM, New York, NY, USA, 45–54. DOI : <http://dx.doi.org/10.1145/2669485.2669500>
40. Roman Rädle, Hans-Christian Jetter, Mario Schreiner, Zhihao Lu, Harald Reiterer, and Yvonne Rogers. 2015. Spatially-aware or Spatially-agnostic?: Elicitation and Evaluation of User-Defined Cross-Device Interactions. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 3913–3922. DOI : <http://dx.doi.org/10.1145/2702123.2702287>
41. Jonathan C. Roberts. 2007. State of the Art: Coordinated Multiple Views in Exploratory Visualization. In *Coordinated and Multiple Views in Exploratory Visualization, 2007. CMV '07. Fifth International Conference on*. 61–71. DOI : <http://dx.doi.org/10.1109/CMV.2007.20>
42. Yvonne Rogers. 2006. Moving on from Weiser's Vision of Calm Computing: engaging UbiComp experiences. 4206 (sep 2006), 404–421. DOI : http://dx.doi.org/10.1007/11853565_24
43. Yvonne Rogers and Sian Lindley. 2004. Collaborating around vertical and horizontal large interactive displays: which way is best? *Interacting with Computers* 16, 6 (Dec. 2004), 1133–1152. DOI : <http://dx.doi.org/10.1016/j.intcom.2004.07.008>
44. Ramik Sadana and John Stasko. 2014. Designing and Implementing an Interactive Scatterplot Visualization for a Tablet Computer. In *Proceedings of the 2014 International Working Conference on Advanced Visual Interfaces (AVI '14)*. ACM, New York, NY, USA, 265–272. DOI : <http://dx.doi.org/10.1145/2598153.2598163>
45. Ramik Sadana and John Stasko. 2016. Designing Multiple Coordinated Visualizations for Tablets. *Computer Graphics Forum* 35, 3 (jun 2016), 261–270. DOI : <http://dx.doi.org/10.1111/cgf.12902>
46. Stephanie Santosa and Daniel Wigdor. 2013. A Field Study of Multi-device Workflows in Distributed Workspaces. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '13)*. ACM, New York, NY, USA, 63–72. DOI : <http://dx.doi.org/10.1145/2493432.2493476>
47. Arne Schmitz, Ming Li, Volker Schönenfeld, and Leif Kobbelt. 2010. Ad-Hoc Multi-Displays for Mobile Interactive Applications. In *The Eurographics Association*. The Eurographics Association, 45–52.
48. Stacey D. Scott, M. Sheelagh T. Carpendale, and Kori M. Inkpen. 2004. Territoriality in Collaborative Tabletop Workspaces. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work (CSCW '04)*. ACM, New York, NY, USA, 294–303. DOI : <http://dx.doi.org/10.1145/1031607.1031655>
49. Frank M. Shipman, III, Catherine C. Marshall, and Thomas P. Moran. 1995. Finding and Using Implicit Structure in Human-organized Spatial Layouts of Information. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '95)*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 346–353. DOI : <http://dx.doi.org/10.1145/223904.223949>
50. Ben Shneiderman. 2002. Promoting Universal Usability with Multi-layer Interface Design. *SIGCAH Comput. Phys. Handicap.* (June 2002), 1–8. DOI : <http://dx.doi.org/10.1145/960201.957206>
51. Anthony Tang, Melanie Tory, Barry Po, Petra Neumann, and Sheelagh Carpendale. 2006. Collaborative Coupling over Tabletop Displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '06)*. ACM, New York, NY, USA, 1181–1190. DOI : <http://dx.doi.org/10.1145/1124772.1124950>
52. Lucia Terrenghi, Aaron Quigley, and Alan Dix. 2009. A Taxonomy for and Analysis of Multi-person-display Ecosystems. *Personal Ubiquitous Comput.* 13, 8 (Nov. 2009), 583–598. DOI : <http://dx.doi.org/10.1007/s00779-009-0244-5>
53. James R. Wallace, Stacey D. Scott, and Carolyn G. MacGregor. 2013. Collaborative Sensemaking on a Digital Tabletop and Personal Tablets: Prioritization, Comparisons, and Tableaux. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 3345–3354. DOI : <http://dx.doi.org/10.1145/2470654.2466458>
54. Mark Weiser. 1991. The computer for the 21st century. *Scientific American* 265, 3 (Sept. 1991), 66–75. DOI : <http://dx.doi.org/10.1145/329124.329126>
55. Robert E. Wood. 1986. Task complexity: Definition of the construct. *Organizational Behavior and Human Decision Processes* 37, 1 (feb 1986), 60–82. DOI : [http://dx.doi.org/10.1016/0749-5978\(86\)90044-0](http://dx.doi.org/10.1016/0749-5978(86)90044-0)
56. Paweł Wozniak, Nitesh Goyal, Przemysław Kucharski, Lars Lischke, Sven Mayer, and Morten Fjeld. 2016. RAMPARTS: Supporting Sensemaking with Spatially-Aware Mobile Interactions. *Proceedings of the SIGCHI Conference on Human Factors in Computing System 2016* (2016). DOI : <http://dx.doi.org/10.1145/2858036.2858491>
57. Johannes Zagermann, Ulrike Pfeil, Roman Rädle, Hans-Christian Jetter, Clemens Klokmose, and Harald Reiterer. 2016. When Tablets Meet Tabletops: The Effect of Tabletop Size on Around-the-Table Collaboration with Personal Tablets. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 5470–5481. DOI : <http://dx.doi.org/10.1145/2858036.2858224>