GIAnT: Visualizing Group Interaction at Large Wall Displays

Ulrich von Zadow, Raimund Dachselt
Interactive Media Lab Dresden
Technische Universität Dresden
Dresden, Germany
{firstname}.{lastname}@tu-dresden.de

ABSTRACT
Large interactive displays are increasingly important and a relevant research topic, and several studies have focused on wall interaction. However, in many cases, thorough user studies currently require time-consuming video analysis and coding. We present the Group Interaction Analysis Toolkit *GIAnT*, which provides a rich set of visualizations supporting investigation of multi-user interaction at large display walls. GIAnT focuses on visualizing time periods, making it possible to gain overview-level insights quickly. The toolkit is designed to be extensible and features several carefully crafted visualizations: A novel timeline visualization shows movement in front of the wall over time, a wall visualization shows interactions on the wall and gaze data, and a floor visualization displays user positions. In addition, GIAnT shows the captured video stream along with basic statistics. We validate our tool by analyzing how it supports investigating major research topics and by practical use in evaluating a cooperative game.

ACM Classification Keywords
H.5.2. Information Interfaces and Presentation (e.g. HCI): User Interfaces. – Graphical user interfaces, Input devices and strategies, Interaction styles.

Author Keywords
collaborative work; multitouch; awareness; territoriality; physical navigation; coupling; visualization; visual analysis

INTRODUCTION
As display prices fall, large vertical interactive displays become more feasible in many settings. Researchers have found them useful for visualization [4], intelligence analysis [2], and developer meetings [7], among others, and there is use in industry as well (e.g., via the MS Surface Hub). Research on wall displays includes a number of studies on wall interaction [3, 4, 6, 19], as well as introducing new interaction concepts and modalities (e.g., [5, 16, 22, 23, 35]). Scientists have investigated different interaction styles, e.g. gestural [10], multimodal [22], and proxemic [5] interaction, as well as a number of interesting research aspects such as locomotion [1, 20], territoriality [3], equality of participation and dominance [9], and collaborative coupling styles [18].

Often, these studies use video analysis and coding (e.g., [6, 18, 19, 30]), a task that remains very time-consuming in spite of coding tools such as VCode [14] and Chronoviz [11]. For example, Jakobsen and Hornbaek [19] manually coded a total of 24:27 hours of video, in addition to using custom analysis software and computer vision. While domain-specific analysis tools in the related field of tabletop collaboration have been developed (e.g., [27, 33]), they are not designed to handle wall display-specific aspects. In particular (and in contrast to tabletop interaction), user movement is an integral part of wall interaction, and visualizing this movement in space and time is thus important for interaction research.

We contribute the Group Interaction Analysis Toolkit *GIAnT* (Figure 1) that aims to fill this gap. Development was driven by our analysis of existing research in collaborative wall display interaction, which in many cases relies on a deep understanding of user positions and movement. From this analysis, we derived a list of specific research topics that would benefit from tool support (referenced as T1-T6 throughout the paper):

- T1: Locomotion. How do users move in front of the large display? How is movement used to access information (e.g., [2, 4, 20])? Under which circumstances do users change positions or switch places [19, 21]? How often and in which situations do users switch between working close to the wall and at overview distance [4]?
- T2: Territoriality [3, 30] and social proxemics [15]. At what social distances from each other do users interact? Is there evidence for territoriality?
- T3: Equality of participation [21, 27]. Are individual users being left out [27] or dominant (e.g., [6, 9])?
- T4: Coupling styles [18]. Are users collaborating closely or working side-by-side?
- T5: Awareness [13, 34] of collaborators' actions and of interface changes. Do the users have a clear mental picture of the state of the interface? Of collaborators' actions?
- T6: User roles. Are there (perhaps application-specific) user roles (e.g., Villains, Micro-managers, Architects [12]; turn-taking or driver/audience [21])?

CHI 2017, May 06 - 11, 2017, Denver, CO, USA

ACM 978-1-4503-4655-9/17/05...$15.00
DOI: http://dx.doi.org/10.1145/3025453.3026006

GIAnT supports analytical studies at large wall displays by providing innovative and focused visualizations that show aspects of multi-user interaction relevant to the above topics. In contrast to video annotation systems that generally only allow analysis of one point in time at once, GIAnT additionally supports working at an overview level, with aggregated data. The toolkit focuses on visualizing time periods of interaction data, allowing users to see information about complete or partial interaction sessions at a glance and thus speeding up analysis significantly. Additionally, GIAnT supports flexible zooms into shorter time periods, supporting seamless transition to a detailed analysis for periods of time where this is needed (details on demand [31]). Further, researchers can switch to video analysis when appropriate as well.

GIAnT is a standalone application and extensible on several levels: It supports adding new data sources, derived data sets, and visualizations. It includes a number of carefully designed visualizations that support work on a diverse set of research topics. We validated GIAnT by analyzing how it supports work on the research topics T1-T6 and by using it to evaluate a cooperative game. GIAnT is freely available under a GPL license[1] and downloadable via github[2].

RELATED WORK

There are a number of related publications that focus on analysis tools for multi-device and multi-user interaction. VisTaCo [33] visualizes touches on a tabletop, using user-coded touches as main data source. CollAid [27] expands upon this by supporting user-specific audio using a microphone array. Marquardt et al.'s Excite [26] supports analysis of proxemics information in conjunction with video coding, in effect automating the video coding process to a degree through queries in a proxemics database. Further, VICPAM [28] supports analysis of interactions with multiple desktop computers, visualizing application-specific data such as window activations on a multi-user timeline. We expand upon these works by supporting interactive wall displays. Significantly, we visualize user movements in addition to interactions, allowing analysis of locomotion (T1), social proxemics (T2) and coupling styles (T4), as well as adapting analysis support to wall displays and corresponding user motions.

Tool support for interaction research also includes a number of video annotation and analysis tools. Among them are Hagedorn et al.'s VCode and VData [14] as well as Burr et al.'s VACA [8], which both focus on traditional video coding, supporting multiple video streams and a timeline with events as well as generic sensor data. Hofmann et al. [17] focus on collaborative annotation in educational settings. Further, Lasecki et al. [24] crowdsource the video coding process, using paid remote workers to code in parallel and thus significantly reducing the time needed. All of these focus on the video coding process and none of them show visualizations that include time periods or aggregate movement and interaction data over time. One work that does show time period visualizations is Chronovis [11], which supports overlays on maps to visualize paths but focuses on airplane pilot interfaces.

[1]https://www.gnu.org/licenses/gpl-3.0.en.html

[2]https://github.org/imldresden/GIAnT

GIAnT SYSTEM DESIGN

GIAnT is a standalone visual analysis application that is based on an extensible architecture. The system design is built around three major layers: 1) Abstract and concrete sources of data, 2) Derived data calculated from the data sources, and 3) Visualizations that can access all data sources and derived data. It works with a rich set of input data, including user position and gaze direction, video and audio feeds, and touches on the wall. From this, derived data items are calculated. These include gaze points on the wall as well as statistics such as the user's movement speed, distance from the wall and touch frequency. Finally, the set of current visualizations includes a timeline view showing user movement and touches, heat map views of wall gaze points and floor positions, and a parallel coordinate plot that shows basic statistics (Figure 1). In the following, we describe the support for data sources and derived data in detail.

Data Sources

In order to incorporate diverse sources of interaction and movement data in a structured fashion, GIAnT works with the concept of *abstract data sources* – types of data that it uses for visualizations. Specifically, we currently support the user position and gaze direction, video recordings, and user-coded touches on the wall as abstract data sources.

Each of these abstract data sources may have one or more *concrete data sources*. Specifically, one source for the user positions can be a motion tracking system, requiring the users to wear instrumented caps (Figure 1d) or similar. Alternatively, cameras could be used to track the users. The users' gaze directions could be measured either through mobile gaze trackers or approximated through the head direction (acquired using a motion tracking system). Further, a per-user visualization of touches requires a source of touches that include user IDs. This can, e.g., be realized by using user-specific tangible markers to open up interaction lenses (Figure 4) or by using a vision-based system that delivers user IDs with touches (CollAid [27] did this for tabletops, and YouTouch! [36] is an option for wall displays). Finally, as long as users do not work in very close proximity to each other, using the user position data to correlate touches with users is a practical alternative.

Derived Data

From the input data, GIAnT calculates a set of derived data and statistics. Currently, this includes gaze points on the wall (calculated from the users' positions and gaze directions), as well as per-user statistics such as the distance travelled, the average distance from the wall, and the number of touches. Additionally, we expect that support for new sets of secondary data will be added for concrete analysis tasks. As an example, detection of F-Formations [25] or touch zones would be possible.

Further possibilities include calculating additional statistics: the average distance to collaborators (T2, T4) and the number of position swaps (T1) are examples. Finally, if awareness of collaborators' actions (T5) is a research topic (and accurate eye tracking data is available), it would make sense to calculate the average number of users who were able to see a touch.

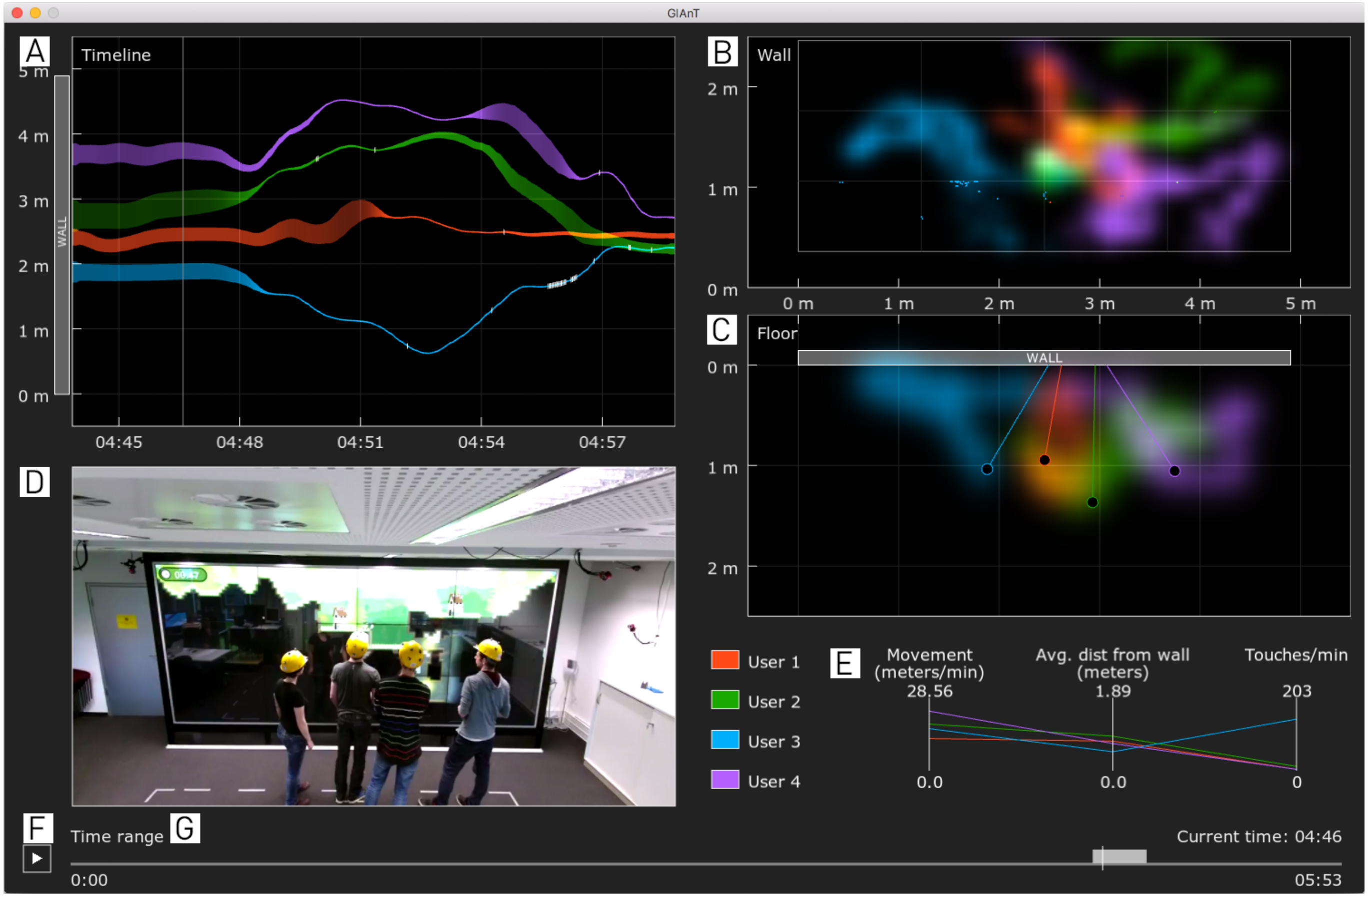

Figure 1. GIAnT user interface showing 15 seconds of interaction: (a) Timeline showing user movement across the wall (distance to wall coded as line width), touches (white dots) and current time (vertical line), (b) Wall visualization showing heat map of gaze points as well as touches (small dots), (c) Floor visualization showing heat map of positions in addition to current user positions (circles) and gaze direction (colored lines), (d) current video frame, (e) parallel coordinate plot of statistics on current time period, (f) play button that toggles realtime playback of the video and statistics, and (g) time slider showing current time range (light grey block) and timepoint (white vertical line).

GIAnT INTERFACE

At the core of GIAnT is a set of visualizations that display a large variety of data in several view panes (Figure 1, (a)-(c)). The challenge here was to visualize user positions, gaze directions and touches, as well as associated changes over time in a way that supports effective analysis of the interaction. In particular, time-dependent user positions are 3D data, which we needed to display on a 2D screen. We thus show user positions in two of the three views: While the Timeline view (a) displays position changes across the wall over time in a line graph, the Floor view (b) displays a heatmap of positions. Additionally, the current point in time as well as time range can be changed interactively. In combination, this allows effective analysis of the dynamics of a situation.

In general, visualizations have access to all of the input data and derived data available to the system, and the set of visualizations is designed to be extensible. The interface also shows a video of the interaction (d) and a statistics pane (e). All data shown is based around the concepts of a current time interval and a current point in time and updated synchronously. The time slider (g) shows the current interval and timepoint, and dragging shifts both. Further, moving the mouse over the timeline visualization changes the current point in time, and the mouse wheel can be used to quickly zoom in and out in the time dimension at this point. Finally, realtime playback is supported by the play button at the lower left of the interface (f). In playback mode, all visualizations are updated synchronously with the video display, giving an integrated view of changes in the users' behavior.

The system shows all participating users by default, but individual users can be shown or hidden using the interface (currently implemented through keyboard commands). Users are color-coded; care was taken to choose colors that are easily distinguishable and have equal perceptual lightness to avoid biases. Therefore, colors are specified in the CIE-LCh color space. Colors have equal L (Lightness) and C (chroma) and are equally spaced along the circular h (hue) axis. To overlay multiple users, we use additive blending, thus side-stepping issues with draw order. As an additional benefit, lightness is left free to visualize other aspects. A minor side-effect of equal lightness is that when converted to greyscale (e.g., for printing), users become indistinguishable. For our current sets of data and on-screen visualization, the results were very satisfactory (see, e.g., Figure 1); however, note that color in visualization is a complex subject [32]. We therefore made extension to other sets of colors easy.

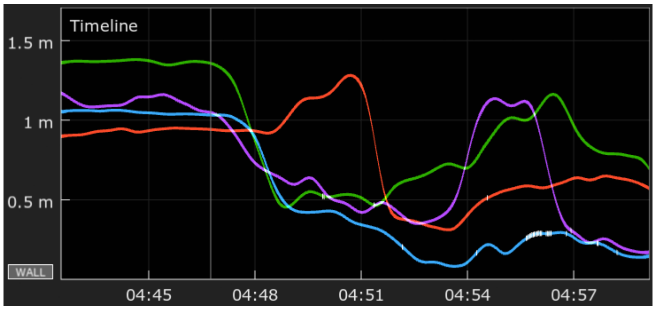

Figure 2. Alternative timeline visualization that shows the distance of the players to the wall, using the same data as in Figure 1. The bottom axis is the current time, the left axis represents the distance.

In the following, we discuss the individual visualizations that GIAnT supports.

Timeline Visualization

Our novel *Timeline Visualization* (Figure 1a) is designed to make the movements of users visible at a glance. The foundation is a per-user line chart with the horizontal axis representing time and the vertical axis representing the position of users across the wall. Additionally, the width of each line is used to show the distance to the wall, with greater width corresponding to greater distance. Conversely, the user's opacity is reduced with distance (see, e.g., Figure 1: at 4:45-4:48, all users stand at overview distance. Later, at about 4:52, all users are close to the wall). The curves are adaptively smoothed to avoid high frequency noise in the data. Further, we visualize touches as white dots superimposed on the lines.

The challenge in designing this visualization was to show three dimensions of data (time and 2D position) in a two-dimensional graph. Dedicating an axis to time focuses the visualization on the dynamics of interaction. In contrast to heat map visualizations, the chronology of events is instantly visible. Further, showing the distance as width allows us to visualize a second spatial dimension, albeit at the cost of a short learning phase. Since width is a visual parameter which is well-suited for quantitative data, the visualization still gives a good indication of the physical distance between users and thus social proxemic zones the users are working at (T2). The visualization also allows good overview of movements (T1) at a glance - users that, e.g., move between overview distance and detail work are clearly visible. Further, differences in movement styles can alert researchers to issues with users being left out (T3) or to different user roles (T6).

In addition to this central view, users can switch to an alternative timeline visualization that uses the vertical axis to directly show the distance of users to the wall (Figure 2). It does not show movement across the wall, but is specialized on visualizing the distance. This is helpful to analyze when and how often users are trying to gain an overview or approach the wall (T1), e.g., in order to touch it. The distance to the wall may also be evidence of dominance (T3), as some tend to stand in the back (e.g., the green user) while others are mostly close to the wall (e.g., blue user). Finally, this visualization helps in analyzing proxemic interaction, where distance is used as an input dimension.

Wall Visualization

The *Wall Visualization* (Figure 1b) shows interactions on the wall. It displays a heat map of gaze points on the wall (whether estimated or measured exactly) as well as a scatterplot of touches. The colors again designate the users, and heat map lightness is mapped to gaze duration. The use of additive blending results in distinctive color changes where several users' gaze points coincide. Further, to allow a clearer view of the touch distribution, the heat map of gaze points can also be hidden, resulting in plots such as the one shown in Figure 3.

Displaying gaze points as heat maps is an established standard in gaze research (e.g., [29]), which we expand to support multiple users by color coding. In our application case, this visualization allows analysts to, e.g., quickly get an approximation of the visibility of touches to other users, an indicator of awareness (T5). Touch locations and clusters are also indicators of on-screen territoriality (T2), with global territories visible when the complete time period is selected, while transient territories become visible for shorter time periods. Further, if users rarely look at the same points on the screen (visible in Figure 1, top right), this may be an indicator of loose coupling (T4).

Floor Visualization

GIAnT's *Floor Visualization* (Figure 1c) shows movement of users in front of the wall from a bird's eye perspective. It overlays a heat map showing positions of users in the current time interval with circles showing user positions and head directions for the current point in time. As before, heat map lightness is mapped to dwell time of users.

As part of our iterative development, we also implemented an experimental alternative floor visualization that used a line graph to show user movement. While this would theoretically have shown changes over time more effectively, it quickly became cluttered when used to visualize longer time periods. Further, it was hard to discern current user position and gaze directions in this variant. The current visualization clearly shows regions that individual users have occupied (e.g., in Figure 1c, the blue user is the only one using the left half of the wall), giving indications of both physical navigation (T1) and territoriality (T2). The degree of overlap in the users' positions is also evidence for the degree of coupling (T4).

Statistics View

In the *Statistics View* (Figure 1e), GIAnT shows a number of statistics generated live for the current time interval as a paral-

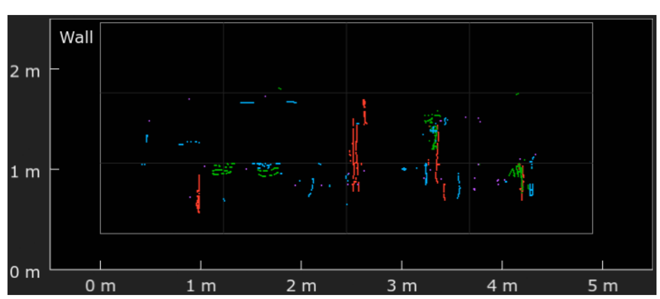

Figure 3. Wall visualization configured to only show the scatterplot of touches on the wall.

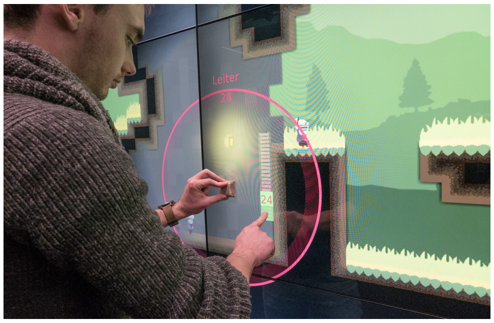

Figure 4. Interaction lens created by using a tangible marker: The user is touching the wall with a tangible marker in his left hand, allowing user-specific interaction with his right hand.

lel coordinate plot. The view currently shows the movement speed, the average distance from the wall and the number of touches per minute; we expect additional derived data to be added to this view as well. We use parallel coordinates for this because they make distributions of data clearly visible (e.g., in Figure 1, bottom right, the blue user is using a lot more touch interactions than the others), and because of their potential to show correlations between attributes and easy extensibility to additional attributes.

The research questions the statistics help answer depend on the values displayed: For example, average distance from the wall and number of touches may be indicators of equality of participation (T3); the average distance from the wall can also show whether the user is looking at an overview or concerned with details (T1).

IMPLEMENTATION

Our application is based on the media and user interface framework libavg[3]; the user interface is scripted in python. To acquire user positions, we currently use an Optitrack[4] motion tracking system, with gaze direction approximated using head tracking. The current system supports generating user IDs for touches by requiring users to place tangible markers to open up user-specific interaction lenses (Figure 4), and expansion to support user IDs through headtracking should be straightforward.

The raw csv-formatted recorded data is preprocessed in a separate step and stored in an SQLite[5] database. This preprocessing step converts device-specific (e.g., Optitrack) data to abstract data (e.g., user position). Hence, the actual application only needs to support abstract data sources. The preprocessor further normalizes time steps and pre-calculates several intermediate values (e.g., low-pass filtered position data) for speed. Additionally, the video needs to be synchronized with tracking and touch data using a precise timestamp. We use a small script which starts the video recording. This script uses a filename that includes the time the recording starts. Alternatively, an initial touch that is visible in the video could also be used to manually synchronize it with the data logs. Further, the videos are converted to a format that does not contain delta frames to allow fast seeking.

[3]https://www.libavg.de

[4]http://optitrack.com

[5]http://www.sqlite.org

An additional implementation goal was to support fluent interaction at a sufficient framerate to allow quick 'scrubbing' through information. Consequently, data needs to be accessed and visualizations generated in realtime. Therefore, significant effort went into optimizations: Speed-critical parts of the application are realized as a libavg plugin written in C++. This includes the rendering of the timeline visualization and heat map, as well as live calculation of statistics. Finally, seeking in videos uses a synchronous (i.e., threadless) decoder for minimal latency.

CASE STUDY: MINERS

In the section on the GIAnT interface above, we not only introduced the visualization views but also analyzed their suitability for the research topics T1-T6. Besides this analytical validation, we validated our toolkit by using it to analyze a cooperative game at an interactive wall display, *Miners* (a corresponding study was published in [34]). Toolkit validation was part of an iterative process: We continually improved our interaction and visualization concepts by using successive prototypes to do practical analysis work on data collected during the study on Miners. To illustrate its utility in a specific application example, we therefore walk through one of our analysis sessions using GIAnT.

A full description of our game is available in [34]; a synopsis follows. Miners is a game similar to Lemmings played by four players at an interactive wall display (image of gameplay in Figure 1d). Players cooperate to rescue miners trapped in a cave. Initially, the cave is completely dark, and each player has a specific tool that manipulates the game world: Two players can build ladders and bridges, respectively. Further, one player can remove obstacles using a pickaxe, and the last has a set of lanterns that light up an area of the cave. Each player has a tangible marker that represents his or her tool. Interaction is bimanual: Players place the marker with one hand, opening up a circular interaction lens. Touching inside the lens activates the tool at that spot (Figure 4). The game forces players to cooperate: They have a shared goal, and levels can generally only be completed by using all tools.

Our case study uses recorded data from a playthrough of a game level of Miners, including head tracking data, video, and complete touch event data from the wall display with user IDs.

Upon opening GIAnT, the tool shows an initial overview spanning the complete session (the timeline view is visible in Figure 5, wall and floor views in Figure 6). The timeline visualization alone allows a number of interesting observations:

- The players generally move as a group and stay close to one another (T4: coupling).
- Player movements show similar patterns and it appears that no player is left out (T3: Equality).
- There are frequent position switches (T1: Locomotion).
- Looking a bit closer, we can tentatively identify four broad movement phases (T1: Locomotion). Phase 1: The first 30 seconds are spent far away from the wall; Phase 2: From this point in time to about 4:30, users generally stay close to the wall, moving as a close group; Phase 3: A phase follows

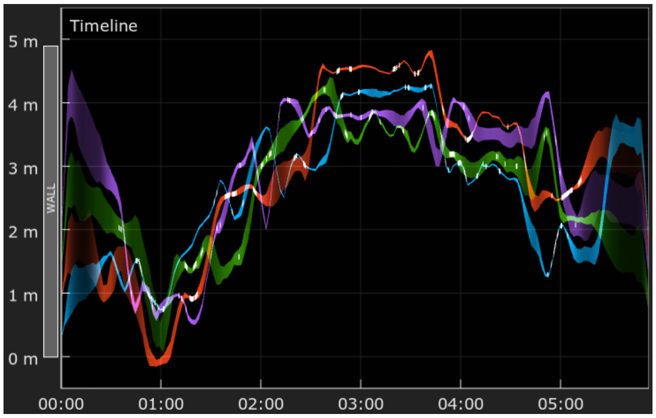

Figure 5. Timeline showing one Miners level playthrough. Among others, it shows at a glance: players move as group, there are frequent position switches, and there is a large time interval (about 1:00-4:30) where players stay close to the wall in general.

in which users disperse and join up again; and Phase 4: From about 5:00, players are again far from the wall.

Further, the wall visualization (Figure 6, top) shows no evidence of private territories (T2), since the touch positions overlap a great deal. Gaze direction is estimated using the head direction in this case. While this means that individual measurements are inaccurate, we assume that the aggregate view shown in the wall visualization approximates reality well. The floor visualization (Figure 6, bottom) shows that almost the complete floor is used by all players, further reinforcing the impression that there are no private territories. The roughly triangular shape also shows that players generally gravitate to the center when at overview distance. Finally, there is a hotspot at the right center of the wall that has received a lot of attention - clearly visible in both heat map views.

All these clues can be gathered in the initial minute of analysis.

To further study the initial session, we then worked to gather more insights at a higher level of detail. We therefore zoomed in so that the views showed 30 seconds of data and moved through the complete session at this granularity. We still found no evidence for territoriality (T2). At this point, we were able to use the video playback to get a clearer understanding of the semantics of the phases identified earlier. Phase 1 is an initial planning phase, and Phase 4 is after the game ends. It turns out that for almost the entire duration of the game (Phase 2), players very seldom move far enough back to see an overview of the complete wall (T1). We also saw physically close interaction, near-constant movement as a group, and similar gaze points throughout.

However, the short Phase 3 (4:30-5:00) seemed to be an exception, because it is the only time at which players do not appear to act as a group. We therefore decided to examine this phase in more detail and zoomed to the time interval from 4:44 to 4:59 (visible in Figure 1). The wall and floor views clearly show that player positions and viewpoints have almost no overlap during this time. Viewing video playback for this interval, we saw the actual gameplay situation: The players are almost done with the level, but there is a single miner left in the dark part of the cave that they need to find. They therefore decide to split up and search for this miner, and gather again when one of them is successful. We concluded that this group cooperated well throughout the complete session.

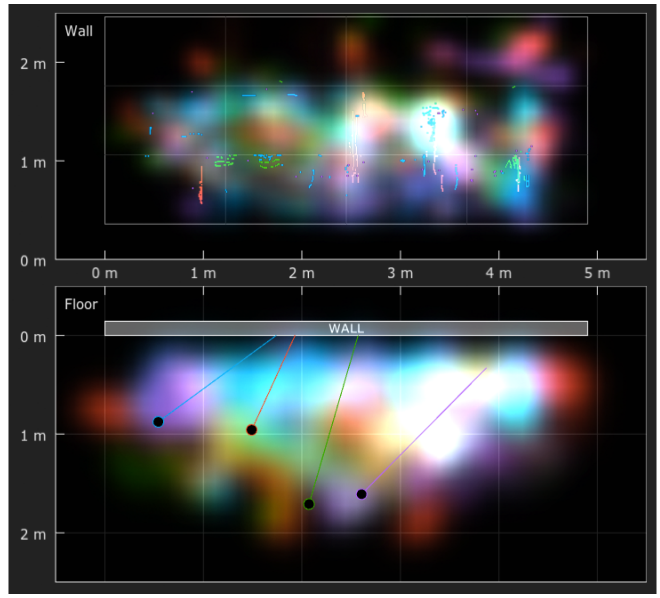

Figure 6. The entire playthrough from Figure 5 in the wall and floor views. Note that touches (dots in the wall visualization) are evenly distributed and show no signs of territoriality. Also visible is that much of the interaction was concentrated at a spot at the center right of the wall.

During this intermediate-level analysis, we also noticed an interesting pattern: There are several situations where a player stays closely behind the others, quickly moves forward to touch, and moves back again in the span of a few seconds. We decided to find out whether this pattern occurs regularly. Both timeline views show this pattern well (Figure 7, green player), and we were able to quickly find other occurrences of the pattern. Again watching the video at these points in time, it became clear that the users in question were moving away quickly to avoid hindering others' interactions. As a follow-up, we could have used additional targeted video coding, e.g., to gather quantitative data.

In total, the analysis had so far taken much less time than the actual running time of the video. Contrast this to analysis using traditional video coding, which would have involved a lengthy process, multiple iterations of coding, and no ability to see an overview at a glance. Instead, we were able to quickly form initial hypotheses involving multiple topics and zoom in to more detailed views to gather evidence on these topics, focusing on particularly interesting sections of the interaction.

CONCLUSION

We presented GIAnT, a toolkit for the analysis of multi-user interaction at large wall displays. Using several carefully-crafted visualizations, GIAnT enables insights on a multitude of research topics, facilitating overview-level analysis at a glance while still supporting detail work when needed. Our toolkit thus significantly speeds up the analysis process. GIAnT was validated analytically (by showing how it supports work on a number of research topics) and by practical use in the evaluation of the cooperative game Miners.

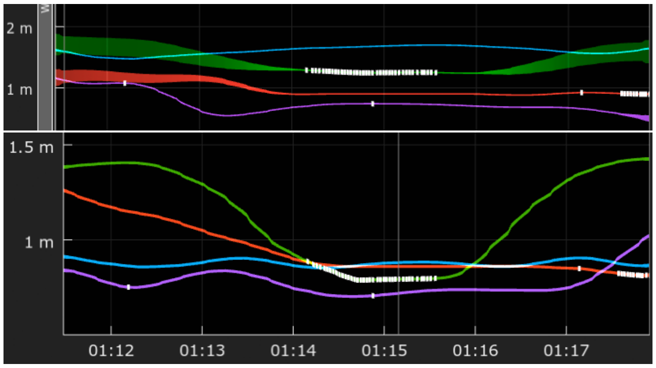

Figure 7. Recurring interaction pattern shown in both variants of the timeline view. While the bottom view very clearly shows movement towards and away from the wall, the top view shows movement across the wall, coding distance from the wall as line width. It is clearly visible that the green user moves towards the wall, touches a number of times (white dots) and quickly moves back a significant distance to make room for other users.

We are currently using GIAnT for our own research, since it allows us to gain a quick overview of interactions with devices and among users. Further, since it is openly available[2], other researchers can benefit from its possibilities as well.

However, it is important to realize that the current tool is a research prototype. As such, it does not implement a full analysis interface – for instance, integrated video coding support is missing, and the data sources currently supported reflect the concrete research task it was used for. Nonetheless, its extensibility should allow us (and others) to easily add support for new data sources, secondary data and visualizations as they are needed for the research tasks at hand. For instance, it is possible to add support for true multi-device and multi-modal (e.g., gestural, pen and tangible interaction) scenarios or to integrate data mining on the recorded interactions.

ACKNOWLEDGMENTS

The authors would like to thank Ricardo Langner and Anke Lehmann for their valuable conceptual input and feedback as well as Alexandra Weiß for initial visualization ideas. David Stolze and Gustav Hahn implemented an initial prototype of the toolkit - thank you. Further, the anonymous reviewers deserve thank for their insightful comments.

REFERENCES

1. Christopher Andrews, Alex Endert, Beth Yost, and Chris North. 2011. Information Visualization on Large, High-resolution Displays: Issues, Challenges, and Opportunities. *Information Visualization* 10, 4 (Oct. 2011), 341–355. DOI: http://dx.doi.org/10.1177/1473871611415997

2. Christopher Andrews and Chris North. 2013. The Impact of Physical Navigation on Spatial Organization for Sensemaking. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2207–2216. DOI: http://dx.doi.org/10.1109/TVCG.2013.205

3. Alec Azad, Jaime Ruiz, Daniel Vogel, Mark Hancock, and Edward Lank. 2012. Territoriality and Behaviour on and Around Large Vertical Publicly-shared Displays. In *Proceedings of the Designing Interactive Systems Conference (DIS '12)*. ACM, New York, NY, USA, 468–477. DOI: http://dx.doi.org/10.1145/2317956.2318025

4. Robert Ball and Chris North. 2008. The Effects of Peripheral Vision and Physical Navigation on Large Scale Visualization. In *Proceedings of Graphics Interface 2008 (GI '08)*. Canadian Information Processing Society, Toronto, Ont., Canada, Canada, 9–16. DOI: http://dx.doi.org/10.20380/GI2008.02

5. Till Ballendat, Nicolai Marquardt, and Saul Greenberg. 2010. Proxemic Interaction: Designing for a Proximity and Orientation-aware Environment. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces (ITS '10)*. ACM, New York, NY, USA, 121–130. DOI: http://dx.doi.org/10.1145/1936652.1936676

6. Jeremy P. Birnholtz, Tovi Grossman, Clarissa Mak, and Ravin Balakrishnan. 2007. An Exploratory Study of Input Configuration and Group Process in a Negotiation Task Using a Large Display. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. ACM, New York, NY, USA, 91–100. DOI: http://dx.doi.org/10.1145/1240624.1240638

7. Andrew Bragdon, Rob DeLine, Ken Hinckley, and Meredith Ringel Morris. 2011. Code Space: Touch + Air Gesture Hybrid Interactions for Supporting Developer Meetings. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces (ITS '11)*. ACM, New York, NY, USA, 212–221. DOI: http://dx.doi.org/10.1145/2076354.2076393

8. Brandon Burr. 2006. VACA: A Tool for Qualitative Video Analysis. In *CHI '06 Extended Abstracts on Human Factors in Computing Systems (CHI EA '06)*. ACM, New York, NY, USA, 622–627. DOI: http://dx.doi.org/10.1145/1125451.1125580

9. Andrew Clayphan, Roberto Martinez-Maldonado, Martin Tomitsch, Susan Atkinson, and Judy Kay. 2016. An In-the-Wild Study of Learning to Brainstorm: Comparing Cards, Tabletops and Wall Displays in the Classroom. *Interacting with Computers* (2016), 1–23. DOI: http://dx.doi.org/10.1093/iwc/iww001

10. Raimund Dachselt and Robert Buchholz. 2008. Throw and Tilt – Seamless Interaction across Devices Using Mobile Phone Gestures. In *Proceedings of the 2nd Workshop on Mobile and Embedded Interactive Systems (Lecture Notes in Informatics Series (LNI))*. Gesellschaft für Informatik e.V. (GI), 272–278.

11. Adam Fouse, Nadir Weibel, Edwin Hutchins, and James D. Hollan. 2011. ChronoViz: A System for Supporting Navigation of Time-coded Data. In *CHI '11 Extended Abstracts on Human Factors in Computing Systems (CHI EA '11)*. ACM, New York, NY, USA, 299–304. DOI: http://dx.doi.org/10.1145/1979742.1979706

12. T.C. Nicholas Graham, Irina Schumann, Mrunal Patel, Quentin Bellay, and Raimund Dachselt. 2013. Villains, Architects and Micro-managers: What Tabula Rasa Teaches Us About Game Orchestration. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 705–714. DOI:http://dx.doi.org/10.1145/2470654.2470754

13. Carl Gutwin and Saul Greenberg. 2002. A Descriptive Framework of Workspace Awareness for Real-Time Groupware. *Computer Supported Collaborative Work* 11, 3 (Nov. 2002), 411–446. DOI: http://dx.doi.org/10.1023/A:1021271517844

14. Joey Hagedorn, Joshua Hailpern, and Karrie G. Karahalios. 2008. VCode and VData: Illustrating a New Framework for Supporting the Video Annotation Workflow. In *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI '08)*. ACM, New York, NY, USA, 317–321. DOI: http://dx.doi.org/10.1145/1385569.1385622

15. Edward T. Hall. 1966. *The Hidden Dimension*. Doubleday.

16. Michael Haller, Jakob Leitner, Thomas Seifried, James R. Wallace, Stacey D. Scott, Christoph Richter, Peter Brandl, Adam Gokcezade, and Seth Hunter. 2010. The NiCE Discussion Room: Integrating Paper and Digital Media to Support Co-Located Group Meetings. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '10)*. ACM, 609–618. DOI: http://dx.doi.org/10.1145/1753326.1753418

17. Cristian Hofmann, Nina Hollender, and Dieter W. Fellner. 2009. *Workflow-Based Architecture for Collaborative Video Annotation*. Springer Berlin Heidelberg, Berlin, Heidelberg, 33–42. DOI: http://dx.doi.org/10.1007/978-3-642-02774-1_4

18. 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 Transactions on Visualization and Computer Graphics* 18, 5 (May 2012), 689–702. DOI:http://dx.doi.org/10.1109/TVCG.2011.287

19. Mikkel R. Jakobsen and Kasper Hornbæk. 2014. Up Close and Personal: Collaborative Work on a High-resolution Multitouch Wall Display. *ACM Transactions on Human-computer Interaction* 21, 2, Article 11 (Feb. 2014), 34 pages. DOI: http://dx.doi.org/10.1145/2576099

20. Mikkel R. Jakobsen and Kasper Hornbæk. 2015. Is Moving Improving?: Some Effects of Locomotion in Wall-Display Interaction. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 4169–4178. DOI: http://dx.doi.org/10.1145/2702123.2702312

21. Mikkel R. Jakobsen and Kasper Hornbæk. 2016. Negotiating for Space?: Collaborative Work Using a Wall Display with Mouse and Touch Input. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 2050–2061. DOI: http://dx.doi.org/10.1145/2858036.2858158

22. Ulrike Kister, Patrick Reipschläger, Fabrice Matulic, and Raimund Dachselt. 2015. BodyLenses: Embodied Magic Lenses and Personal Territories for Wall Displays. In *Proceedings of the 2015 International Conference on Interactive Tabletops & Surfaces (ITS '15)*. ACM, New York, NY, USA, 117–126. DOI: http://dx.doi.org/10.1145/2817721.2817726

23. Ricardo Langner, Ulrich von Zadow, Tom Horak, Annett Mitschick, and Raimund Dachselt. 2016. *Content Sharing Between Spatially-Aware Mobile Phones and Large Vertical Displays Supporting Collaborative Work*. Springer International Publishing. DOI: http://dx.doi.org/10.1007/978-3-319-45853-3_5

24. Walter S. Lasecki, Mitchell Gordon, Danai Koutra, Malte F. Jung, Steven P. Dow, and Jeffrey P. Bigham. 2014. Glance: Rapidly Coding Behavioral Video with the Crowd. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology (UIST '14)*. ACM, New York, NY, USA, 551–562. DOI: http://dx.doi.org/10.1145/2642918.2647367

25. 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

26. Nicolai Marquardt, Frederico Schardong, and Anthony Tang. 2015. EXCITE: EXploring Collaborative Interaction in Tracked Environments. In *Proceedings of INTERACT 2015*. Lecture Notes in Computer Science, Vol. 9297. Springer International Publishing, 89–97. DOI: http://dx.doi.org/10.1007/978-3-319-22668-2_8

27. Roberto Martínez, Anthony Collins, Judy Kay, and Kalina Yacef. 2011. Who Did What? Who Said That?: Collaid: An Environment for Capturing Traces of Collaborative Learning at the Tabletop. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces (ITS '11)*. ACM, New York, NY, USA, 172–181. DOI: http://dx.doi.org/10.1145/2076354.2076387

28. Roshanak Zilouchian Moghaddam and Brian Bailey. 2011. VICPAM: A Visualization Tool for Examining Interaction Data in Multiple Display Environments. In *Proceedings of the 2011 International Conference on Human Interface and the Management of Information - Volume Part I (HI'11)*. Springer-Verlag, Berlin, Heidelberg, 278–287. http://dl.acm.org/citation.cfm?id=2027916.2027951

29. Jakob Nielsen and Kara Pernice. 2009. *Eyetracking Web Usability* (1st ed.). New Riders Publishing, Thousand Oaks, CA, USA.

30. Stacey D. Scott, S. Carpendale, and M. Inkpen. 2004. Territoriality in collaborative tabletop workspaces. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work (CSCW '04)*. ACM, 294–303. DOI: http://dx.doi.org/10.1145/1031607.1031655

31. Ben Shneiderman. 1996. The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations. In *Proceedings of the 1996 IEEE Symposium on Visual Languages (VL '96)*. IEEE Computer Society, Washington, DC, USA, 336–343. DOI: http://dx.doi.org/10.1109/VL.1996.545307

32. Samuel Silva, Beatriz Sousa Santos, and Joaquim Madeira. 2011. Using color in visualization: A survey. *Computers & Graphics* 35, 2 (2011), 320 – 333. DOI: http://dx.doi.org/10.1016/j.cag.2010.11.015

33. Anthony Tang, Michel Pahud, Sheelagh Carpendale, and Bill Buxton. 2010. VisTACO: Visualizing Tabletop Collaboration. In *ACM International Conference on Interactive Tabletops and Surfaces (ITS '10)*. ACM, New York, NY, USA, 29–38. DOI: http://dx.doi.org/10.1145/1936652.1936659

34. Ulrich von Zadow, Daniel Bösel, Duc Dung Dam, Anke Lehmann, Patrick Reipschläger, and Raimund Dachselt. 2016. Miners: Communication and Awareness in Collaborative Gaming at an Interactive Display Wall. In *Proceedings of the 2016 ACM on Interactive Surfaces and Spaces (ISS '16)*. ACM, New York, NY, USA, 235–240. DOI:http://dx.doi.org/10.1145/2992154.2992174

35. Ulrich von Zadow, Wolfgang Büschel, Ricardo Langner, and Raimund Dachselt. 2014. SleeD: Using a Sleeve Display to Interact with Touch-sensitive Display Walls. In *Proceedings of the Ninth ACM International Conference on Interactive Tabletops and Surfaces (ITS '14)*. ACM, New York, NY, USA, 129–138. DOI: http://dx.doi.org/10.1145/2669485.2669507

36. Ulrich von Zadow, Patrick Reipschläger, Daniel Bösel, Anita Sellent, and Raimund Dachselt. 2016. YouTouch! Low-Cost User Identification at an Interactive Display Wall. In *Proceedings of the International Working Conference on Advanced Visual Interfaces (AVI '16)*. ACM, New York, NY, USA, 144–151. DOI: http://dx.doi.org/10.1145/2909132.2909258