Beyond Post-It: Structured Multimedia Annotations for Collaborative VEs

João Guerreiro[1], Daniel Medeiros[1,2], Daniel Mendes[1], Maurício Sousa[1], Joaquim Jorge[1], Alberto Raposo[3], Ismael Santos[4]

[1]INESC-ID / Instituto Superior Técnico / Universidade de Lisboa, Lisboa, Portugal
[2]CAPES Foundation, Ministry of Education of Brazil, Brasilia, DF, Brazil
[3]Tecgraf Institute of Technology, PUC-Rio, Rio de Janeiro, Brazil
[4]CENPES, Petrobras Research Center, Rio de Janeiro, Brazil

Figure 1: *Our approach enables creating multimedia annotations and provides a tree-based structure to support traceability in decision making. (a) An annotation in the VE, created by user "Obama". It also shows how to create a new annotation. (b) Annotation as it is playing. It relies on camera movements (synchronized with audio) to show an anomaly in the model. (c) The discussion structure. Balloons indicate annotations' type and colors denote arguments' type (Pro, Con, +Info). The user is adding a (fourth) positive argument (Pro) with an audio annotation, as it can be seen by active buttons.*

Abstract
Globalization has transformed engineering design into a world-wide endeavor pursued by geographically distributed specialist teams. Widespread adoption of VR for design and the need to act and place marks directly on the objects of discussion in design reviewing tasks led to research on annotations in virtual collaborative environments. However, conventional approaches have yet to progress beyond the yellow postit + text metaphor. Indeed, multimedia such as audio, sketches, video and animations afford greater expressiveness which could be put to good use in collaborative environments. Furthermore, individual annotations fail to capture both the rationale and flow of discussion which are key to understanding project design decisions. One exemplar instance is offshore engineering projects that normally engage geographically distributed highly-specialized engineering teams and require both improved productivity, due to project costs and the need to reducing risks when reviewing designs of deep-water oil & gas platforms. In this paper, we present an approach to rich, structured multimedia annotations to support the discussion and decision making in design reviewing tasks. Furthermore, our approach supports issue-based argumentation to reveal provenance of design decisions to better support the workflow in engineering projects. While this is an initial exploration of the solution space, examples show greater support of collaborative design review over traditional approaches.

1. Introduction

The engineering design process usually involves teams of specialists from different disciplines and globalization has transformed it in a geographically distributed collaborative environment. The adoption of Collaborative Virtual Environments (CVEs) for design reviewing enables parallel work and improved efficiency [Hua02]. Moreover, it supports collaborative meetings that focus and act on the actual virtual models, instead of communication based on voice and video alone. The need to place marks directly on the objects of discussion in design reviewing tasks led to research on annotations in CVEs.

Several research projects contributed with textual or sketch annotations placed on the virtual world, following the post-it + text metaphor (e.g. [JdAB*10, JGD02, UCFM10]). Such annotations enable users to identify problems and point out their suggestions/opinions in context. However, conventional approaches lack the expressiveness and efficiency of co-located meetings where participants may resort to speech, sketching on a paper or object manipulations to illustrate their opinions. We believe that multimedia annotations can provide similar expressiveness through audio, sketch and by reproducing camera movements in synchrony with speech and sketch.

Annotations serve to influence, guide and document project decisions. This makes them important for subsequent collaborative sessions. Still, managing the knowledge exchanged is one of the major challenges in CVEs [LTA09]. Most approaches present annotations as isolated and unrelated items, which undermines both structured discussions and decision making. A central tenet of our approach relies on provenance, which "*pertains to the derivation history of a data product starting from its original sources*" [SPG05], to reveal the design decisions' flow and support discussion in engineering projects.

One exemplar domain that could leverage multimedia annotations is offshore engineering. Projects involving deepwater oil & gas platforms can take several years to complete and engage geographically distributed highly specialized teams. Moreover, they are expensive and have major risks involved [dSSR12]. Multimedia annotations could assist design review meetings in CVEs both by enriching the participants' arguments and increasing expressiveness. Furthermore, providing a tree structure to such annotations/arguments would reveal the provenance of design decisions and empower participants in the discussion.

In this paper, we present an approach based on rich, structured multimedia annotations to support the decision making process in design review tasks. The multimedia alternatives intend to ease creating annotations (sketch and audio; text for titles) and augment them with synchronized camera movements in the virtual environment. Moreover, we rely on a tree-based structure and issue-based argumentation both to support discussion and to reveal the provenance of design

decisions. In such processes, one may respond to a particular annotation with additional information or contribute with positive/negative arguments.

2. Related Work

Annotations are important to augment a collaborative environment with additional information, such as reporting errors or suggesting corrections. A mainstream example is *Microsoft Word* and its *track changes* feature. In a collaborative VR environment, the literature reports a predominance for post-it-like textual and/or sketch annotations (e.g. [dSSR12, JdAB*10, JGD02, UCFM10]). Although less frequent, audio annotations have also been used, for example to tag the surroundings in a Mobile Augmented Reality Platform [LRZS13].

A different way to support discussions is through virtual guided tours [dSSR12, KDTCP08], where a user may share his view of the model with other users. In such settings, one may create virtual paths through critical spots in the scene to demonstrate or discuss possible anomalies. Although this feature is valuable in synchronous meetings, it falls short of supporting asynchronous interaction. Yet, such a feature can be encapsulated in an annotation. In fact, previous works show that recording an avata, r's movement and audio is able to support persistent virtual environments, with a sense of co-location even though the recordings were previously made [IJL*99, LJDB99].

Although useful to identify problems with the model, isolated annotations do not support a discussion. A conventional chat tries to provide such support [UCFM10], but its use is mainly limited to synchronous discussions. Lenne et al. [LTA09] present a valuable contribution by adding semantics to annotations. Their approach enables filtering annotations related to a specific concept and adds different types of arguments (e.g. reacting to a proposal). Yet, it is not clear how different annotations are related and what was the discussion workflow. In contrast, related work in collaborative environments focusing on the semantic grid, use a clear structure that aims to support the before, after and actual meetings [BBSCB*04, DRFMP06]. Design Reviewing tasks in CVEs, in both synchronous and asynchronous forms, could take advantage of a similar structure to support the entire decision making process as well.

3. Design Rationale

Current approaches fall short on providing a framework for asynchronous decision making in CVEs design review. Herein, we present the design rationale that supported our approach to structured multimedia annotations. Our context involves work developed on CEDAR, a Collaborative Virtual Environment for Design Review in the Oil Industry. In what follows, we will briefly describe CEDAR to provide adequate context, discuss our case for multimedia annotations

in order to motivate our approach to enriching multimedia (mainly non-textual) annotations with explicit support for argumentation and collaborative discussions.

3.1. CEDAR System

CEDAR proposes a design review tool to support collaborative tasks from the oil industry field. We followed a user-centered methodology starting with a task analysis involving engineers operating traditional design review CAD tools on a daily basis. Via both structured interviews and informal talks, we were able to clearly identify both challenges to be tackled and real usage scenarios. In oil industry settings, large teams usually gather to review, manipulate and discuss around large CAD models, which are often difficult to visualize and to navigate owing to their complexity. These teams typically include field engineers working at remote locations and technical staff operation at a central location. Maintenance in specific infrastructures inside an oil platform is very frequent and technical staff have to plan how to access them beforehand. Furthermore, changes to these installations are frequent, removing and installing new components. One key goal of ours is to support engineers' field work by allowing them to visualizing in loco the procedures needed to repair devices while assessing possible interferences that are not reflected in the 3D model but exist in the real scene in close cooperation with headquarters.

Our system was developed and evaluated using two geographically-distinct facilities based using large scale displays with different specifications. One is a 2x2 tiled display system using four projectors in rear-projection mode located in Europe as depicted in Figure 2. The mosaic generates 3.6 megapixel images and the visualization system is driven by a single workstation behaving as a single desktop and thus simplifying the development and deployment of applications. The second visualization facility is located in South America. It is composed of four projectors in a trapezoidal arrangement creating a cave-like system as shown


Figure 3: *CEDAR Powerwall in South America*

by Figure 3. While lateral screens use rear-projection, the floor uses front-projection. The projectors are connected to a workstation through, which uses the mosaic feature and each screen is oriented side-by-side. Unlike traditional cave-type systems, we prefer a 60 degree lateral glass screen configuration, allowing either simultaneous visualizations by a sizable audience or several users to operate within the environment.

3.2. System Architecture

While a detailed description of the architecture is beyond the scope of this paper, we offer a synopsis here. The CEDAR
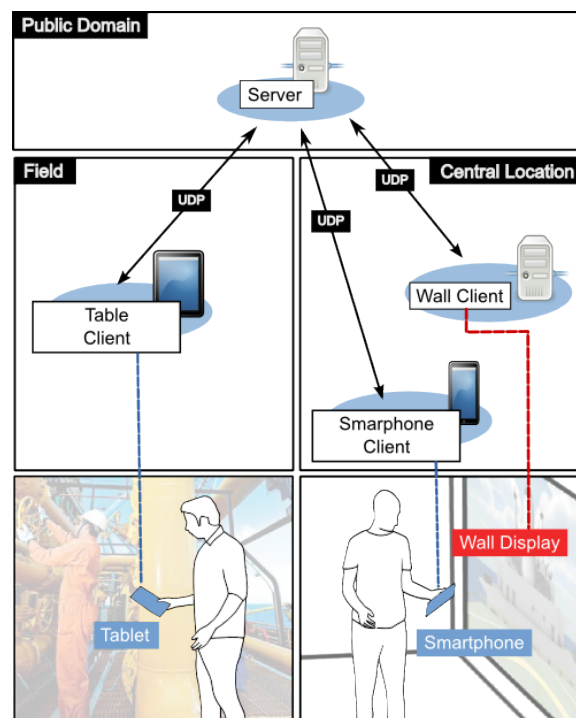

Figure 2: *CEDAR Powerwall in Europe*

Figure 4: *CEDAR system's architecture.*

system follows a star network topology, allowing different clients spread across the globe to connect to a central server and communicate among themselves, as depicted by Figure 4. We consider two different clients to interact with the CAD model under revision during a review session. A client running on a desktop computer connected to a large scale display provides a unique visualization environment for several users discussing engineering problems. Other client, running on an iPad or Android handheld device, offers a personal window displaying its own view of the scene. This view affords multi-touch 3D navigation and can be synchronized (or not) to other clients, namely the large screen display. The server coordinates the different clients and manages the engineering review session, by forwarding communications. By using the client application on each visualization facility, users can connect to the server using an Internet connection and start sharing data between locations in a transparent manner.

3.3. Collaborative Features

Besides providing 3D model visualization and navigation, the CEDAR system supports several collaborative features. The following aim to improve the collaboration between users by providing both direct and indirect communication, visibility and awareness.

Speech Communication: This allows remote, natural communication between the users, which is crucial to the success of the collaborative tasks;

Viewport Sharing: The point of view presented on a display can be shared with other users, allowing geographically apart teams to see the same. Suggested uses are: Displaying the location of small objects, guided tours and simply knowing were the other colleagues are;

Avatars: An avatar representing the current position and orientation is presented on the environment. This affords faster recognition of where other users are located and their look–at direction;

Target Identification: When users point to specific locations in the scene, the system highlights target objects at that position in addition to displaying an avatar. Thus, other users can know which object is being referred to, without requiring verbal descriptions;

Annotations: During reviewing sessions, it is often necessary to take notes. Instead of taking these on a notebook or using a different application, CEDAR offers both sketched (using the handheld device) and speech annotations, that can be viewed or replayed later by users. These are extended in the present work.

3.4. Exploring the Virtual Environment

Our CVE enables users to navigate in the virtual environment to visualize the model using different devices, as it supports touch and mouse navigation. Both touch and mouse approaches for navigation rely on ThumbCam [MSFJ14],

offering look around, move around and scrutinize actions using a single contact point. The camera can be rotated by placing and moving a touch and two rotation modes are supported: look around and circle around (or scrutinize). Two sequential taps on a point in the virtual environment toggle between these two modes. Touching and holding the screen will apply the Drag'n Go technique [MMG12] to move forward. Upwards, downwards and sideways movements can be achieved by tapping and immediately touching the screen, i.e. a quick sequence of touch begin, touch end and another touch begin actions.

While navigating, users may create annotations just by entering scrutinize mode, by focusing on the object which should be annotated, and selecting the annotation menu. This annotation becomes available to all users in the virtual environment persistently, represented by a speech balloon icon. A click or tap on the balloon opens a specific menu presented in Figure 6. Then, the original annotation can be played back as any of the replies. Moreover, it is possible to reply to any node, by selecting both the answer (pro, con, additional information) and annotation type (audio, camera-path, sketch). While interacting with the CVE it is possible to answer to such annotations. Moreover, users joining the CVE at a later time will also have access (and the opportunity to reply) to the annotations. These interactions directly support asynchronous meetings and discussions.

3.5. Multimedia Annotations

Collaborative design review requires support for discussion among participants. Such design review sessions usually come together with notes and minutes that point out current problems, solutions and next steps. When performing design review using a CVE the discussion may be centred on the actual model, whereas the participants may annotate specific parts or identify issues therein.

Annotations in a VE are often restricted to text-based post-its and/or sketches placed on the virtual world. Our multimedia annotations support text-based titles, sketches and audio annotations. Moreover, we enable creating dynamic annotations that include camera movements synchronized with both sketch and audio. Such synchronized annotations resemble video as they can reproduce all steps performed by the author. The rationale behind these annotation types lie on the fastness and easiness to create annotations and on the power they provide in order to point out a problem, solution or opinion. In particular, sketches (more illustrative) and audio are faster to create and are richer than textual notes. However, textual information is also important to present information beforehand or to resume a set of annotations - so we use them on annotation titles.

Camera movements enable a participant to guide others throughout the model or specific parts of it. Synchronized with audio and sketch in a multimedia annotation, they af-

ford powerful ways to expose a project's problems and provide suggestions. The asynchronous nature of the annotations make possible that any user can revisit this guided tour and add more information or make questions about it. For example, one may navigate in the virtual model while talking about the problems found in particular objects, and suggest alternatives via sketch, or keep using solely audio.

3.6. Structured Annotations Reveal Provenance

Annotations can support a decision making process and therefore are important for subsequent collaborative sessions. However, isolated annotations do not readily accommodate such processes and conventional chats hamper parallel discussions as they rely on a single thread. In a chat-like discussion it becomes very difficult to keep track of parallel discussions or embedded arguments. Such problems are augmented if we consider that the design reviewing process in a CVE can often be asynchronous.

In order to support decision making, each annotation added should become part of a structured argument and provide a solid background to the entire discussion. This way it could provide an understanding of both ongoing arguments to identify open issues, help in their resolution and reveal the process used to solve a particular problem in closed issues. Such background is akin to provenance which is used in many contexts (e.g. location and ownership chronology of a work of art; or scientific research reproducibility). For instance, *Vistrails* captures the provenance of a visualization (how it was created), as well as the data manipulated [SFC07]. This enables both reproducibility and a greater understanding of the entire process, which may be of use in similar visualizations.

Similarly to [LTA09], we rely on metadata (e.g. author, position, time) to support accountability for individual annotations. Most importantly, we propose a tree-based structure that both supports discussion within annotations and discloses the provenance of design decisions. Our structured tree is inspired on issue-based models used to capture argument structure, support decision-making and allow subsequent understanding of the decisions taken (Issue-Based Information Systems [KR70, BYC90]).

In our issue-based discussion, any participant may open a new issue, which may include a question or pinpointing a problem with the engineering model under review. Moreover, this annotation may also contain a suggested solution. Afterwards, participants may respond to a particular annotation by adding additional information such as questions, details or taking a stand, or adding positive/negative arguments (Pros and Cons) to the discussion. The ability to answer any of these annotations, from root down to the leaves, creates the aforementioned tree. This both affords a clear identification of the argument flow and supports parallel discussions about an open issue which may emerge.

As a result of this clearly identified argument structure, we are able to support the whole CSCW matrix [Bae95], as it augments the annotation scaffold to allow asynchronous collaboration.

4. Designing Structured Annotations

Our approach to structured multimedia annotations was built on top of CEDAR, the previously described multi-platform CVE. This collaborative system, based on Unity3D, allows using multiple devices such as laptops, wall displays, tabletops and handheld devices, such as tablets or smartphones. Such diversity enables both individual interactions as well as team interactions, either co-located or remote. For instance, participants may share their tablet's screen with the wall, or with all other participants' devices, while manipulating the virtual model, in order to explain their point of view. Figure 5-b portrays another example, where an user is interacting with the discussion tree in his *smartphone*, while using the wall to present the actual annotation for all users in the room to see.

Design review tasks demand greater flexibility when teams are geographically distributed. Still, the main challenges concern asynchronous iterations that may arise from the lack of availability of members in highly-specialized teams, due to timezone differences and offline work. Our approach overcomes that problem by providing a flexible framework to support both discussion and decision making.

4.1. Annotations in the Virtual World

As noted earlier, our system supports multimedia annotations through text (for a small description) sketch, audio and synchronized camera movements (evocative of, but more useful than a video annotation). Static sketches are bi-dimensional drawings entered via mouse or finger touches on the screen and mapped to the near plane of the main camera. The software maintains the camera static during the operation, to ensure a consistent view of sketched strokes. Sketches can also be synchronized with camera movements and audio. In this manner users can enter multiple sketches, each associated to a camera position. Both sketches and camera trajectories are saved on an external file indexed by the note. Audio content is handled separately and saved in a different file which can be reproduced simultaneously with sketches and camera movements. These annotations are easily created in the virtual world through a menu that is always visible at the bottom, which shows the possible annotation types when activated, as depicted in Figure 1-(a).

Annotations are represented in the virtual world using comics speech balloons, that metaphorically signify a discussion. To avoid cluttering the virtual world in popular discussion threads, the number of balloons drawn is the base 2 logarithm of all annotations. This can indicate the relative activity (number of annotations) of an issue's discussion, but

Figure 5: *(a) presents the oil platform model used in this scenario. (b) presents an user interacting with the annotations tree, while the wall presents only the model*

liberates the virtual world of additional clutter to allow visualizing and interacting with the engineering model under review.

We relied on graphical elements to portray an annotations' current state and relations between elements (or nodes) in a discussion. For instance, the border color of an annotation in the virtual world indicates whether that issue has been already solved (green) or not (red). Moreover, the issue author's profile picture is presented in the center of the speech balloon (Figure 1-a). A click/tap on the balloon displays the entire discussion contained in that annotation and enables the current user to participate in the discussion,by adding argumentative nodes.

4.2. Tree-Based Discussion

After selecting an annotation in the virtual world, the corresponding menu is presented on the left side of the screen (Figure 1-b,c). One may select one of the annotations and play it (and then pause it) by selecting the button at the bottom. All annotations are presented in the exact same conditions as they were created (same position and camera orientation).

Every participant may reply to a specific annotation by adding Pro, Con or Additional information nodes (e.g. question or show support) and selecting the annotation content (Figure 1-c). After completing the annotation, a new node is created indicating its content (audio, sketch or camera movements). Its edge color indicates the response type (**Pro** in green; **Con** in red; **+Info** in yellow). Such information is

useful to provide an overview of the discussion workflow without the need to play all annotations. In order to avoid cluttering this view, the tree is re-arranged such as to become centered on the selected annotation. For instance, Figure 1-b depicts a collapsed node (+) on the leftmost side, indicating more activity therein.

Finally, users with permissions allowing them to close such issues, may declare them solved (or re-open them as needed).

4.3. Scenario

We explore our approach via a real user scenario to illustrate all the functionalities developed. The scenario covers a design review engaging geographically distributed engineers. These users can record their interactions on the virtual environment, leave notes and audio annotations about the issues considered crucial, both for the construction of the real model and the virtual (CAD) model of the platform (Figure 5-(a)).

Figure 6 depicts the annotations in the virtual world and discussion workflow, based on a real scenario that occurred in a design review:

1. John asks whether the valve should not have a mechanism to indicate if it is closed or open.
2. Sally, the Project Manager replies that this module does not work that way and this is the usual valve they use. It closes clockwise. She suggests keeping it.
3. Mark adds a positive argument, mentioning that workers are used to this valve and know how to work with it.

4. Sally adds that they have several of these valves in stock, which is another advantage to using them.
5. John insists that such indication is important, mainly for recent workers that may not be aware of how it works.
6. Sally states that it is included in their training, so that is not a problem.
7. John agrees to use that valve then
8. Sally realizes the valve is in the wrong position and uses camera trajectories to show it. She asks the designer to correct it ASAP.
9. When the problem is worked out, Sally marks the issue as solved.

5. Discussion and Conclusions

Design review in CVEs requires a flexible framework to support discussion. Annotations can ease the intra-team communications by placing marks directly on the object of interest. However, conventional annotations lack both flexibility and expressiveness to visualize and support multi-threaded
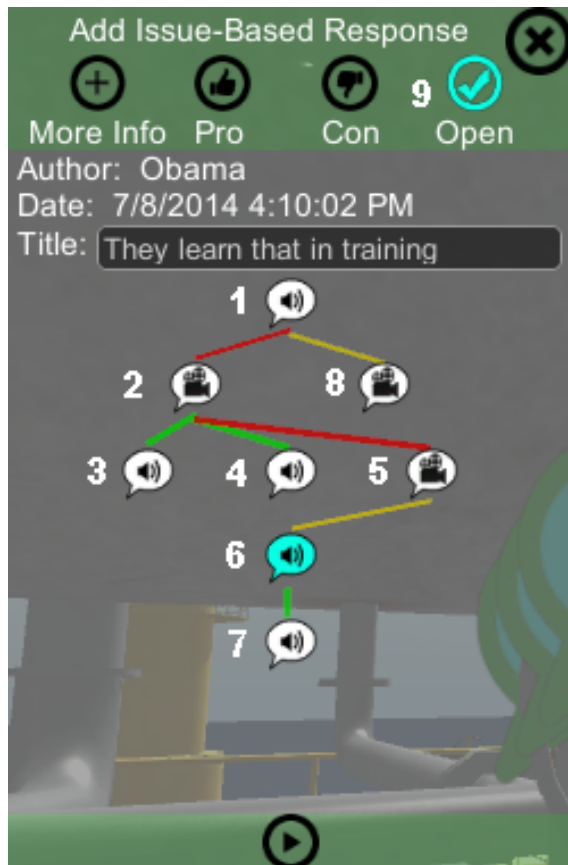

Figure 6: *An example of the discussion tree, based on a real design reviewing scenario. The numbers indicate the order of events.*

discussions and complex arguments. We provide distinct multimedia options to maximize expressiveness, where a video-like annotation is the emblematic example. One can play a sketch or audio annotation in context, where viewing conditions are replicated to show the authors' perspective. Still, it is with camera movements (position and orientation) annotations that users can better explain and demonstrate their points of view. As seen in our scenario, one may manipulate the camera while talking in order to show an anomaly that seems non-existent, but only becomes visible from a specific perspective. The prototype was iteratively developed in close collaboration with an oil company's research team that gave us constant feedback on improvements and requirements to enhance our approach. Moreover, we focus on real use cases that company workers have faced to address their needs, which yielded the scenario presented in this paper. We also focus on the design rationale and principal constituents of our prototype, while presenting a fleshed out example that we played in virtual meetings.

Our multimedia annotations afford complete and expressive arguments, yet unstructured media do not support discussion by themselves. Indeed, our scenario discussion would not be supported by isolated, unrelated and unstructured annotations. Our issue-based structure allows users to reply to specific arguments (annotations) in a discussion with well-identified arguments (Pros and Cons). Such a structure is very important to support decision making and to reveal provenance of design decisions. In our scenario, the process that involved the oil platform valves can be only fully understood by playing the related annotations in ways that a clutter of postit-line annotations could not.

While this paper showcases a virtual model in an advanced design phase with lots of detail, another scenario that could be addressed is the conceptual (early) design of a virtual model in similar engineering projects. Still, additional efforts are needed to supporting highlighting and presenting the changes that occur in a 3D virtual (CAD) model over time. As future work we are also studying ways to better visualizing and navigating through the annotations on the virtual environment. Our approach could also be extended to support creating training- and maintenance-oriented content in easy and natural ways by end–users.

6. Acknowledgements

References

[Bae95] BAECKER R. M.: *Readings in Human-Computer Interaction: toward the year 2000*. Morgan Kaufmann Pub, 1995. 5

[BBSCB*04] BACHLER M., BUCKINGHAM SHUM S., CHEN-BURGER J., DALTON J., DE ROURE D., EISENSTADT M., KOMZAK J., MICHAELIDES D., PAGE K., POTTER S., ET AL.: Collaborative tools in the semantic grid. 2

[BYC90] BURGESS YAKEMOVIC K. C., CONKLIN E. J.: Report on a development project use of an issue-based information system. In *Proceedings of the 1990 ACM Conference on Computer-supported Cooperative Work* (New York, NY, USA, 1990), CSCW '90, ACM, pp. 105–118. URL: http://doi.acm.org/10.1145/99332.99347, doi:10.1145/99332.99347. 5

[DRFMP06] DE ROURE D., FREY J., MICHAELIDES D., PAGE K.: The collaborative semantic grid. In *Collaborative Technologies and Systems, 2006. CTS 2006. International Symposium on* (2006), IEEE, pp. 411–418. 2

[dSSR12] DOS SANTOS I. H., SOARES L. P., RAPOSO A.: A collaborative virtual reality oil & gas workflow. *International Journal of Virtual Reality 11*, 1 (2012). 2

[Hua02] HUANG G. Q.: Web-based support for collaborative product design review. *Computers in Industry 48*, 1 (2002), 71–88. 2

[IJL*99] IMAI T., JOHNSON A. E., LEIGH J., PAPE D. E., DE-FANTI T. A.: The virtual mail system. In *Proceedings of the IEEE Virtual Reality* (Washington, DC, USA, 1999), VR '99, IEEE Computer Society, pp. 78–. URL: http://dl.acm.org/citation.cfm?id=554230.835724. 2

[JdAB*10] JOTA R., DE ARAÚJO B. R., BRUNO L. C., PEREIRA J. M., JORGE J. A.: Immiview: a multi-user solution for design review in real-time. *Journal of Real-Time Image Processing 5*, 2 (2010), 91–107. 2

[JGD02] JUNG T., GROSS M. D., DO E. Y.-L.: Annotating and sketching on 3d web models. In *Proceedings of the 7th international conference on Intelligent user interfaces* (2002), ACM, pp. 95–102. 2

[KDTCP08] KLEINERMANN F., DE TROYER O., CREELLE C., PELLENS B.: Adding semantic annotations, navigation paths and tour guides to existing virtual environments. In *Virtual Systems and Multimedia* (2008), Springer, pp. 100–111. 2

[KR70] KUNZ W., RITTEL H. W.: *Issues as elements of information systems*, vol. 131. 1970. 5

[LJDB99] LEIGH J., JOHNSON A. E., DEFANTI T. A., BROWN M.: A review of tele-immersive applications in the cave research network. In *Virtual Reality Conference, IEEE* (1999), IEEE Computer Society, pp. 180–180. 2

[LRZS13] LANGLOTZ T., REGENBRECHT H., ZOLLMANN S., SCHMALSTIEG D.: Audio stickies: visually-guided spatial audio annotations on a mobile augmented reality platform. In *Proceedings of the 25th Australian Computer-Human Interaction Conference: Augmentation, Application, Innovation, Collaboration* (2013), ACM, pp. 545–554. 2

[LTA09] LENNE D., THOUVENIN I., AUBRY S.: Supporting design with 3d-annotations in a collaborative virtual environment. *Research in Engineering Design 20*, 3 (2009), 149–155. 2, 5

[MMG12] MOERMAN C., MARCHAL D., GRISONI L.: Drag'n go: Simple and fast navigation in virtual environment. In *3D User Interfaces (3DUI), 2012 IEEE Symposium on* (March 2012), pp. 15–18. doi:10.1109/3DUI.2012.6184178. 4

[MSFJ14] MENDES D., SOUSA M., FERREIRA A., JORGE J.: Thumbcam: Returning to single touch interactions to explore 3d virtual environments. In *2014 ACM International Conference on Interactive Tabletops and Surfaces* (2014), ACM. 4

[SFC07] SILVA C. T., FREIRE J., CALLAHAN S. P.: Provenance for visualizations: Reproducibility and beyond. *Computing in Science & Engineering 9*, 5 (2007), 82–89. 5

[SPG05] SIMMHAN Y. L., PLALE B., GANNON D.: A survey of data provenance in e-science. *ACM Sigmod Record 34*, 3 (2005), 31–36. 2

[UCFM10] UVA A. E., CRISTIANO S., FIORENTINO M., MONNO G.: Distributed design review using tangible augmented technical drawings. *Computer-Aided Design 42*, 5 (2010), 364–372. 2