

Collaborative Augmented Reality Environments: Integrating VR, Working Materials, and Distributed Work Spaces

Monika Büscher², Michael Christensen¹, Kaj Grønbaek¹, Peter Krogh³,
Preben Mogensen¹, Dan Shapiro², Peter Ørbæk¹

¹Department of Computer Science
University of Aarhus
DK-8200 Aarhus C, Denmark
{toby, kgronbak, pmogensen,
poe}@daimi.au.dk

²Department of Sociology
Lancaster University
Lancaster LA1 4YL, UK
{m.buscher, d.shapiro}
@lancaster.ac.uk

³Aarhus School of Architecture
Danish Center for Integrated Design
DK-8000 Århus C
peter.krogh@a-aarhus.dk

ABSTRACT

This paper describes *Manufaktur*, a prototype of a concept and infrastructure that goes beyond the classical CVE systems toward a collaborative augmented reality environment, where users' documents and objects appear as live representations in a 3D workspace. *Manufaktur* supports collaborative, distributed work across platforms. It provides the possibility of display on the MS Windows and the SGI IRIX platforms with mono and stereographic displays ranging from a monitor built into a table to a Holobench.

The emphasis is on representing the *materials* with which people are working, and the actions they are taking with them, and far less on the visual representation, e.g. through avatars, of the collaborating *persons*. It seeks to facilitate users' collective self-organisation of environments, spaces and traces that embody the trajectory of their work.

Keywords

Virtual Reality, Working Material, Virtual Office/Project Room, Roomware, CSCW, 3D Workspace.

INTRODUCTION

A number of collaborative virtual environments have been developed in recent years. DIVE [5, 12] and MASSIVE [8] are among the most well known. Both of these systems provide a combination of 3D worlds, avatars, video, sound and documents. Users are represented as simple avatars and they can communicate by means of various media. For instance, avatars may walk up and share a Whiteboard where they can manipulate 2D graphics, watch live video, write and draw simultaneously. They can also import 3D objects, which can be inspected from arbitrary positions. Documents are implemented as portable whiteboards that users can bring with them and be manipulated privately by a single user. Documents can be placed on shared whiteboards or conference tables. A conference table is a service to allow group discussion. A user joins the discussion when his/her avatar gets sufficiently close to the table.

CVE researchers have experimented with different ways of mediating awareness and formation of groups and meetings in virtual environments. However, the artifacts related to the work of users have received less attention in this research, e.g., by not supporting native document formats.

The documents that most CVEs support are low quality writing on the Whiteboard, i.e. manipulation of real office documents and objects is not supported. This is a serious barrier for bringing CVEs out of entertainment or research [1] to real world work tasks. Such integration of documents and objects is a main design issue in the work of this paper.

The work draws on experience and techniques from a number of different fields: Collaborative Virtual Environments (CVE) [6], Open Hypermedia [19], and CSCW [14, 24]. In previous papers [3, 18] we described an early version of a system, the *Manufaktur* ('craft workshop'), intended to provide a computer-based environment to support the work primarily of architects, landscape architects and other designers. This paper outlines a number of further developments to this system, some of which are well on the way to realisation and others of which are in process.

Design professionals use a great variety of materials – plans, sketches, diagrams, photographs, scale models, samples of materials, catalogues, etc. – and the system is based on observations of the ways in which they constantly arrange, relate and rearrange these in three dimensions in their physical workspaces. We found that manipulating the presence and absence of materials, bringing them into dynamic spatial relations, and referring between them, are not just a context or prerequisite for doing the work; rather, they are an integral part of accomplishing the work itself. What is important is not just to create or change a document or other materials, but to do so in the presence of and in relation to others.

The environment is being designed on the basis of the *Manufaktur* system, which was originally designed as an MS Windows 3D environment, integrating documents and design models in common workspaces for (landscape) architects. *Manufaktur* is now becoming a multi-platform distributed concept for augmented reality work environments, where design models and document objects (document representations) can be visualized either on a Windows platform or on a Silicon Graphics platform, on mono or stereographic displays ranging from a monitor built into a table to a Holobench. Document objects and models are, independently of their rendering and display, linked to the real documents and model sources, which may reside on

arbitrary servers on an Intranet, an Extranet or the Internet. The environment supports distributed collaboration in a variety of loosely and tightly coupled modes. Compared to earlier CVEs the environment provides collaboration support that focuses on the working materials. For instance, users' mutual awareness is supported through the material being worked on instead of showing an avatar. Tightly coupled collaboration is supported through synchronous sharing of workspaces and even the camera with which the users see the workspace. The software infrastructure is maturing and we are working on setting up rooms to be tested by users from the domain of (landscape) architecture and industrial design. Among the important design issues discussed in the paper are display and interaction techniques, awareness support, portability of workspaces, and bridging gaps between physical objects in the room and their computer representations.

MANUFAKTUR

The *Manufaktur* provides a digital environment that supports relations between working materials. It supports the configuration of spatio-temporal order by allowing people to construct familiar configurations of documents and objects in a 3D virtual workspace. The same thing can be the work object for one activity, background material or context for another activity, and would be an irrelevant distraction for a third – all in ways that change on a moment-to-moment basis. The *Manufaktur* supports this through the situated creation and manipulation of context, awareness, and action. Its basic features, presented in detail in [3, 18], are summarised below.

Manufaktur Workspace

Figure 1 shows a screen dump from an early prototype of the *Manufaktur*. It shows an open workspace containing (references to) a set of objects. Double clicking any of the document objects will launch it in its respective application with that document, and changes to it will be updated within 3D *Manufaktur* in near real time. The objects can be sized, moved, rotated, etc; light effects may be applied; documents can be made (semi) transparent; organized into groups; and much more.

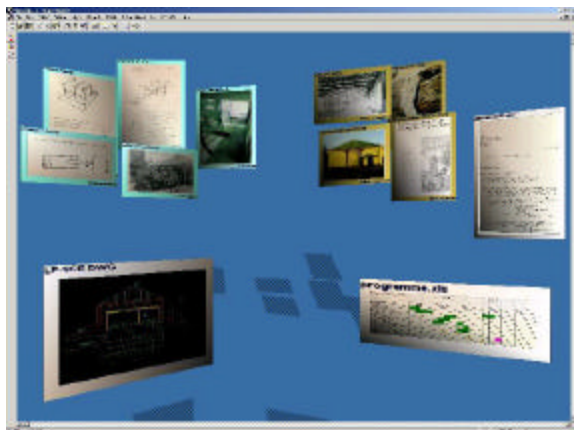


Figure 1: Interface of *Manufaktur*

The *Manufaktur* handles workspaces, which are sets of

references to objects in a certain spatial order. Thus the same object may be present in many workspaces, i.e. the same object is part of various contexts of work. Several different types of objects can populate the space:

Document-objects - 'Live' documents from third party 2D applications with an ActiveX document server, such as spreadsheets, drawing programs, CAD programs, etc. The documents are displayed on objects in a 3D rendered OLE/ActiveX container.

3D models – Various 3D models may be inserted into a workspace, as artefacts worked, objects providing topology to the space, iconic reminders, etc.

'Implantations' – objects or devices that support the customising of a space to changing uses, for example, to create spatial partitions or for imprinting specifically expressive codes [15]. Objects to constrain space, semi-transparent objects to define areas, lights to indicate particularities, shadows on the 'ground' to provide a sense of distance are all examples of 'implantations'.

Groupings – representations of sets of objects (including groups) that may be manipulated as a whole. Besides the implicit grouping of placing objects in proximity, groups can be explicitly defined and manipulated.

Endpoints – representation of link anchors on document-objects. By integrating and extending the Webvise [9, 10] hypermedia system, we enable linking capability to parts of document-objects residing within *Manufaktur*.

A 3D environment like this not only makes it possible to have many documents visible at the same time, but it also, by spatial proximity, makes it possible to indicate the (changing) relevance of a document for work-in-progress, and its relation to other objects in the space. This means that there are in effect different 'levels of openness' of a document, which can still be identified from far away. We see this as a possibility for supporting fluent relationships, and the flexibility to zoom into a detail and out to see the whole, to simultaneously hold present a large number of parameters and their relationships.

Client architecture on the Windows Platform

The *Manufaktur* application on the Microsoft Windows 2000 platform builds on the Microsoft Foundation Class (MFC) framework. Most importantly, this framework provides us with a host of templates and wrapper classes for working with ActiveX/OLE. Furthermore it provides a number of classes for automating and encapsulating work with databases. For rendering 3D scenes we utilise the MS DirectX programming interface.

As already mentioned, documents pertaining to applications external to the *Manufaktur* can be attached to 3D boxes inside a workspace. The specific technology involved in communicating with these external applications is called OLE linking. When one inserts documents into *Manufaktur* (e.g. via drag-and-drop from the file explorer) we establish an OLE link to the file, hereby subscribing to update events for the file. When the external application

"fires" an event, we ask it to draw the document into an off-screen device context (the applications are 2D and so cannot draw themselves directly as 3D objects in *Manufaktur*). We grab the content of the device context and wrap it onto a surface on one of our 3D objects.

The persistent version of these "live" OLE links is called a moniker. Through a so-called "binding" process, monikers are resolved to the documents they refer to, hereby making a running version of them available (i.e. re-establishing the OLE link). Because we are using the distributed version of COM, DCOM, these monikers can be resolved to documents residing anywhere on a network, as long as the machine supports DCOM.

The software architecture of *Manufaktur* is based on the MFC document/view architecture. The document part, corresponding to a workspace, holds the data, and the views handle display of the data. A number of classes provide an object-oriented model of the workspace's data. In this sense it both provides a model of the objects that can actually be rendered to a view and also a model of what is saved to the database. A view of a workspace is a scene containing 3D rendered objects. The rendering is handled by the Direct3D and DirectDraw parts of the DirectX SDK. Direct3D consists of two separate parts, a low-level interface called Immediate Mode and a high-level interface called Retained Mode. Careful use of the low-level interface may enhance performance. We have, however, chosen to use the high level Retained Mode interface mainly due to expected speed-up in development and its support for hierarchies. On the Windows platform, Direct3D/DirectDraw was chosen over the obvious alternative OpenGL because the COM-based DirectX was expected to be better suited for mixing with ActiveX/OLE – technologies also based on COM. The advantages have not been as obvious as expected, so we may experiment with OpenGL, and also with Direct3D Immediate Mode to gain performance.

The above is directed towards acquiring 'live' representations of users' work material. This is in some respects a radical approach to collaborative virtual environments. The emphasis is far less on the visual representation, e.g. through avatars, of the collaborating *persons* – though of course that can also be realised if necessary – than on the representation of the *materials* with which they are working, and the actions they are taking with them. This can be accompanied by speech channels [23] and other affordances of activity, such as 'pointing' by identifiable participants.

WORKSPACES IN PRACTICE

The first level of the *Manufaktur*, described above, enables users to shape, populate and inhabit a three-dimensional digital workspace that replicates some of the features of their physical workspace. By comparison with a standard 2-dimensional windows interface, very many more recognisable documents and other objects can be present with their contents visible or partly visible, without the space appearing cluttered. And this permits the placing and moving of

materials to express the relationships between them, as in the physical workspace.

However, the physical workspace also imposes constraints, and often forces materials into relationship where this is a hindrance not a help. The interplay of these qualities is illustrated in the following brief summary of a passage in the work of a studio of landscape architects.



Our starting point coincides with an early morning chat between two landscape architects. Vic is snowed under, while all of Steve's projects are waiting for decisions elsewhere. He offers to help Vic. A drawing related to work on a park (*Northpark*) is pinned up between two lights on the ceiling – there was no room for it on Vic's drawing board, so he extended his workspace to be able to see it. Steve has worked on this project himself before and inspects the drawing with interest to see what is currently being done.

But Vic is also working on several planting plans for the redesign of an urban road (*Westroad*). The main areas of planting need to be summarised in an overview. This is a drawing Steve can assist with. A first step is to define a number of 'viewports' in a blank AutoCAD file. To get started, Steve needs an idea of how many viewports are required and how they should be arranged. He turns to his drawing board, where he had been looking at materials related to the design of a different park (*Southpark*). To consider possible ways of providing access to a small stream in this park, he had examined a set of technical drawings of steps copied from reference books, the design brief, and some photographs of the existing condition of the park. Now he puts the photographs back into their folder. The detail drawings are pushed to the top left, on top of a folder with timesheets. Steve begins to draw a sketch layout. Vic picks up one of the plans from his drawing board, places it over Steve's documents, and they draw up a plan.

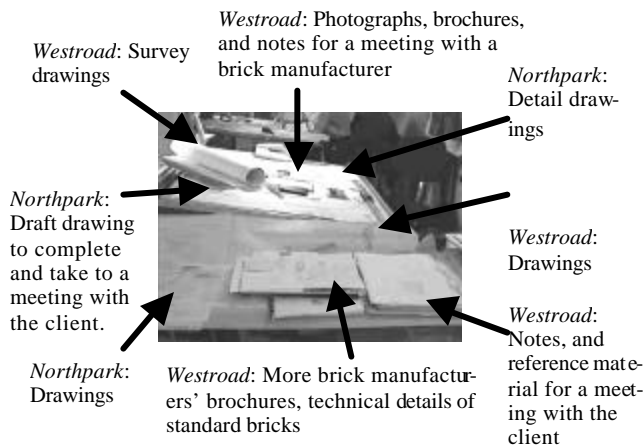


This short excerpt of interaction highlights several features of the role of material arrangements in the physical workspace:

- A publicly visible remnant of a busy day, like the drawing pinned up between the lights, displays some aspects of ongoing work. Material assemblies or parts of them can be a resource for colleagues' mutual awareness of work and the pressures others are under.
- There is a flexible division of labour – colleagues can be drawn into ongoing work as necessary and as possible. New contextual arrangements of documents and objects can be created on the spur of the moment.

- Steve's workspace already contains a set of contextual material arrangements that relate to his projects. They recede into the background as this small new contextual arrangement is formed.
- Steve's arrangements have a spatial order that is disturbed by this activity. The photographs that bear a relation to the design of access to the water are now mixed in with others and have lost their link to technical details of possible solutions.

The notion of spatial order can be developed when we take a look at Vic's workspace a day later:



Vic is working on two projects in parallel, preparing for meetings in the afternoon. The first thing he does in the morning is to assemble the documents he needs for these meetings. On postit notes he scribbles some reminders of things he needs to check before he goes. They stick out like flags from the pages of the documents they relate to, or are pinned to the top of the pile. His main task, however, is to complete a draft of earth movements and site clearance for *Northpark*.



He walks across to shelves with product information and returns with a catalogue. The height of kerbs around a pavilion needs to be determined. Leafing through the catalogue, he also spots a detail drawing

of 'slit drainage', which he might need later that week when producing a whole set of contract drawings for *Northpark*. He marks the page with a box of staples and continues to look for the kerbs. He cannot find the right page and remembers that he took it to a different meeting and left it at home. Nevertheless, he places the folder on a desk behind him, open at the page showing the 'slit drainage'. Both the technical details for the kerbs and the 'slit drainage' will be needed for the final set of contract drawings.

A client phones and reports a problem relating to a different project. Vic assures him that the drawing he sent is correct, promises to check and get back to him in writing. He goes to his machine and browses through the set of CAD draw-

ings produced for this job. There is also some correspondence with the land surveyors updating the data this drawing is based upon. Vic fetches the job-folder to make sure that he has indeed used the correct information. Leaving the folder next to his machine he then returns to his work on *Northpark*.

This second snapshot provides a glimpse of the kind of work involved in considering and relating different parameters that impinge on a design and the traces this work leaves in the material environment. The carefully considered relationships between a diversity of documents at the beginning of the workday results in groups and spatial arrangements that correspond with the temporal order of tasks and their urgency. It is complemented by links made as an integral part of doing the work. They, too, precipitate in spatial arrangements such as the link between specific pages in the catalogue and the draft on the table. In the context of current practices, such material arrangements are ephemeral. They are easily disturbed by other work carried out in the same space and dissolved when a job is done. This is an important resource lost.

Projects can last from several months to several years and over the course of a project different landscape architects contribute to the work or take control of particular sub-tasks. Multiple projects are active at any one time, and in their day-to-day work the landscape architects often divide their attention between them. Every time one switches focus, returns to a project, or is handed responsibility for an unfamiliar project the relationships between a diversity of materials have to be re-established.



While Vic is preparing for his meetings, he remains available for Steve who needs answers to some questions. As Steve stops by, Vic pulls up one of the *Westroad* drawings, places it on top of his *Northpark*

material, and turns it so Steve can see it. In word and gesture they determine how different parts of the planting plans could be turned and scaled to fit into the overview Steve has created. But Vic wants to postpone final decisions to the next day, because there may be changes that arise at his meeting with the client today. This frees Steve to consider some samples sent in by manufacturers for the design of an urban square he and his senior colleague Lynn are in charge of.



They lay out the samples on the floor. The issue they are particularly concerned with is the intersection between three different materials. The central semi circle in front of the town hall entrance is to be laid with expensive natural stone setts, while, towards the outer edge, a stone

similar in colour, texture, and style but less costly has to be found. Both areas are to be broken up by lines of small dark paving stones. While the samples for the central part of the square, and those for the dark lines are accepted as suitable, the sample they thought might serve to pave the outer perimeter turns out to be mismatched. They discuss the qualities a suitable candidate should have with constant reference to the assembly of stones on the floor, other samples, and a sketch of the design displayed on the wall. Then they begin to draw in some product information catalogues.



This third example draws attention to the intertwining of talk, gesture and embodied action, physical objects, and paper copies of documents. Reference and judgements are made in a complex field of textual, visual, and material information. Product information is available in electronic formats – either on CDROM or via the Internet. But a small display on a screen does not dovetail well with the way in which this information needs to be drawn into the work. To address this issue, the boundaries between electronic environments and the physical world need to be made more permeable.

As well as the creative configuration of physical spaces, these examples also show the problematic and messy ways, familiar to us all, in which materials crowd in on each other, often creating an ‘archaeology’ of layers of unrelated materials. This problem can partly be avoided in the *Manufaktur* as a result of two features. First, views can be saved and later opened, so that there can be multiple workspaces, each targeted at a particular project or sub-project. These can be closed or minimised and later returned to, either to cope with a short-term interruption or a longer-term hiatus in the life of a project. Second, the *Manufaktur* contains a reference onto an underlying document or object, not a copy of it. This means that the same item can have a presence in many different workspaces without giving rise to version problems.

The consequence is that users can, in effect, inhabit as many ‘parallel universes’ in relation to their work as they care to create. This constitutes a second level of use of the *Manufaktur*. It is a very powerful feature, but could also be very confusing and disorientating in use. An important aim for the fieldwork will be to discover how users choose to subdivide their digital environments, and to propose aids for navigating among them, and for recording and representing historical traces of their paths.

These snapshots on the one hand illustrate some of the rationale for the first prototypes of *Manufaktur* – Multiple workspaces with spatially arranged document-objects. On the other hand they point to needs for expanding the basic

notion as well. In what follows, we describe what has been accomplished in relation to two of these avenues: Collaboration, and augmenting the physical environment.

COLLABORATION

The early prototype considers single workspaces and multiple workspaces, but for single users. However, the fieldwork extract above, and the thrust of all our research, serves to emphasise the shared and collective character of the work carried out in this kind of community of practice [2, 16]. If the *Manufaktur* is usefully to replicate aspects of the workspace, it must do so as a collaborative environment as well as a virtual one, and this constitutes the third level of its design and use. The *Manufaktur* has an integration and distribution architecture which enables various forms of collaboration to be realised in distributed project groups via Intranets, Extranets or the Internet.

Distribution architecture.

The *Manufaktur* supports real-time distributed interactive collaboration through its network distribution architecture. Geometry and textures of 3D objects are stored in one or more shared relational database servers. Collaborating *Manufaktur* clients interact via a collaboration server used to distribute messages about object updates among session participants.

A *Manufaktur* client enters a collaborative session by subscribing to events from one or more workspaces on the collaboration server. Whenever the client modifies a database object, it notifies the collaboration server about what was changed. Object data is not distributed through the collaboration server; it is shared via the database server.

However, while moving around 3D objects in real-time, coordinates are not committed to the database until the object is released. Instead, object coordinates are distributed through the collaboration protocol, as this is vastly more efficient than going through the database.

As each *Manufaktur* client could be working on several workspaces on several databases at once, all collaboration messages contain the name and port of the database server, the database and table name, and the user name of the sender. Sending this information in each collaboration message while interacting in real-time would require several hundred bytes per message. At frame rates of, say, 60 Hz a simple movement of an object could require a data stream from a single client of as much as 100 kbps. The collaboration server would then replicate this data stream to all the participants in the session.

To improve interactive response times and scalability the collaboration protocol uses differential messaging to compress protocol messages. Collaboration messages come in two flavours: reference messages and difference messages. The collaboration server uniquely identifies each client, and each client keeps track of the last reference message from each participant in the session. Difference messages contain differences from the last reference message sent, and receiving clients use their saved reference messages to recon-

struct whole messages from difference messages. This technique can cut down the size of difference messages by a factor of five.

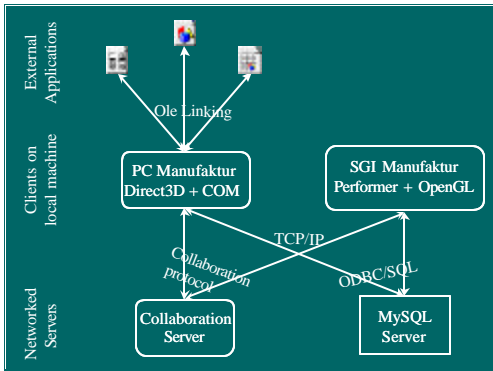


Figure 2: Architecture

Update events contain a bitmask identifying the fields of the corresponding database record that have been updated. Allowing the other collaborators to fetch only these fields from the record speeds up collaboration, as large textures do not have to be reloaded unless they have been updated.

As described, a central collaboration server distributes events among session participants on top of TCP/IP. The choice of TCP-based distribution originates from its ease of implementation and assured message delivery, as well as the lack of a widely deployed IP multicast infrastructure. Using IP multicast (as done in MASSIVE-2 [7]) to distribute events would be more efficient than going through a central server, but is left for future work.

As we saw in the example from the landscape architects, design work is collaborative and deeply intertwined with the handling of documents and objects. Increasingly, there is a need for support of distributed work. Telework, collaboration between colleagues located in different branches, but also collaboration with professional partners (e.g. architects, engineers, surveyors), as well as a need to draw on other people when out on site are becoming more and more commonplace. Furthermore, the landscape architects participating in our research envisage a use of the *Manufaktur* as a tool for presenting ongoing work to their clients. This raises issues pertaining to distribution, various platforms (ranging from a fully equipped (virtual) project room, over clients running on workstations, to a ‘viewer’ running inside a WWW browser), and access rights (among people from the same branch, same company, other company, etc.). Figure 2 depicts the implemented distribution architecture.

Currently we support an individual (uncoupled) mode and a tightly coupled collaboration mode [13, 26], where the virtual project rooms are synchronized and all users are aware of ongoing activities in the viewable area of the room. In this setting we see a natural distinction between two types of tight coupling. The first one is a mode where all object manipulations are reflected in all clients in the same session, and each user’s actions are visible and distinguishable, but the users may place themselves (their camera) as

they wish. The second tightly coupled mode is WISIWYS (What I See Is What You See) where some users share the same camera in the workspace, which they are all allowed to move. The first of the two tight coupling modes is primarily intended for “awareness” purposes, whereas the second is primarily intended for synchronous collaboration. The latter ensures that multiple participants can draw others’ attention to particular points in the visual field by looking at them. It would have to be extended through pointing support.

There are various ways in which collaboration through this infrastructure could be supported technically and through the development of virtual forms of collective practice. The dynamic configuration of collaboration spaces could be supported by allowing users to bring *Manufaktur* views or zones into contact. One could envisage a mixed space in which, for example, the left zone was used by one participant, the right zone was used by another participant, and the middle zone was used by them for material worked on jointly. Other, more peripheral participants might ‘overview’ the whole of this space. As was the case with the ‘parallel universes’ of multiple workspaces, an important aim for the fieldwork will be to discover how users choose to configure spaces for collaboration, given such open possibilities, and how to support them in making these spaces intelligible.

In general, the software architecture, as outlined above, allows for experimentation along (*at least*) four dimensions with regards to collaboration issues:

1. Sharing of views on workspaces/objects within workspaces.
2. Granularity of updates when document contents and/or objects within workspaces are changed or added.
3. Access rights to workspaces/objects within workspaces.
4. Access rights to document files referenced from objects in *Manufaktur* workspaces.

Applying these to the concrete work examples above could result in a scenario like the following: The snowed under Vic might be giving Steve a quick overview of his current projects and they discuss what specifically needs attention by browsing through *Northpark* and *Westroad* workspaces. Initially they share the view. However, after narrowing down the specific task that Steve is supposed to assist with (planting overview for *Westroad*), Vic instantiates a new camera (view) for Steve, guiding him through the documents relevant to the task. Alternatively or later on, a new shared “planting overview” workspace is instantiated and Vic, at the same time having an overview of several other relevant workspaces, throws in the relevant documents from these other workspaces, currently not visible to Steve. On initiating the task Steve has full access rights to change and add to the new workspace, but only has read-access to Vic’s existing workspaces pertaining to *Westroad*. Furthermore, on file level, Steve of course has write-access to the new AutoCAD drawing, but only read-access to some

of the other documents, brought into the workspace from existing *Westroad* workspaces. To focus on the specific task at hand, he decides not to have these other documents updated if Vic is making changes to them concurrently. On the other hand, after making a change Vic thinks it relevant to Steve's task and chooses to notify him. Steve, having chosen not to see document updates in real-time, is in some way discreetly made aware of Vic's change and hereafter explicitly chooses to refresh the changed document object.

The above example illustrates the potential in the various combinations of sharing, update granularity, and access rights. However, it is obvious that a careful design is needed in order not to force the user into spending more time making explicit decisions about the combinations than actually accomplishing work.

AUGMENTING PHYSICAL ENVIRONMENTS

While we are seeking to replicate some features of the physical environment in the *Manufaktur*, we make no assumption that a digital environment is better or that users ought to migrate to it. Rather, we recognise that professional designers *already* do a large part of their work in a digital environment, and that it provides them only poor support. While the digital environment will always have disadvantages, it has some advantages too. It may overcome some of the boundedness and inflexibility of physical settings, and it can contain active objects as well as passive ones. Our purpose is therefore to maximise the capacities of digital environments, so as to maximise choice for users of how to 'flow' their activities between digital and physical worlds. This requires also maximising the capacity to 'bridge' in both directions, through such means as fast printing and scanning, and tagging real objects with links to their electronic representatives, or providing for their automatic recognition. This constitutes a fourth level of the design and use of the *Manufaktur*.

The issue of bridging between the physical and digital worlds has implications for the form that the realisation of the 3D digital environment will take. The simplest, of course, is as 2.5 dimensions on a monitor. But our intention is to explore all of the current, developing and future means of monographic and stereographic representation. This is not really a 'fifth level' of the design and use of the *Manufaktur*, because it is orthogonal to the other four, all of which could be realised in 3D in a variety of ways. But it has particular consequences when the fourth level, involving maximised bridging between the digital and the physical is brought into consideration. For these purposes, a semi-immersive stereographic environment offers the most promise. Stereography will greatly enhance the intelligibility and usefulness of the space. Semi-immersion, e.g. with shutter-glasses, will allow fluid boundaries between the virtual environment and the still-visible real-world environment.

Client architecture on the SGI platform.

For these reasons we wish to bring the *Manufaktur* into the world of (semi-)immersive VR, involving the exploitation

of stereo projection, head tracking and high bandwidth 3D interaction techniques. As our main platform for this is a Holobench backed by a Silicon Graphics Onyx2-IR2 computer, an implementation of the *Manufaktur* for the SGI platform has been carried out.

The *Manufaktur* client on the SGI platform builds on top of the IRIS Performer C++ class library that provides a layer on top of OpenGL corresponding to the retained mode of Direct3D. It also supports the use of multi-channel stereo displays as present on the Holobench, and it provides loaders for a number of third party 3D graphics models.

The software architecture of the SGI client is divided into three layers: an abstract database layer, a distributed storage layer and a 3D object layer. The database layer abstracts the details of the particular SQL database used to hold the shared workspaces and their objects. It also interfaces to the dynamic texture server used to generate textures for external documents and images on the fly.

The distributed storage layer interfaces between the 3D object layer and the database layer and takes care of the interaction with the collaboration server. The 3D object layer contains the representation of the 3D objects (such as document containers, billboards, groups, menus, and virtual tools), and it interacts with IRIS Performer to maintain the scene graph used for display.

Navigation as well as interaction with the objects in a workspace is supported via special 3D objects called tools. Tools range from ordinary keyboard and mouse tools supporting a 2.5D interface on the SGI platform to true immersive tools like a ray-casting tool for 3D picking, a flying tool for navigation and a world-in-miniature tool supporting both picking and navigation. All the tools share a common event distribution infrastructure, and may delegate events further among each other.

Using the Holobench.

Initially, the *Manufaktur* client for the SGI platform is intended for use at a Holobench display device from TAN [28]. The Holobench consists of two back projected 2m × 1m screens mounted at a 90 degree angle. Using head tracked stereo display on both screens enables virtual 3D models to appear on the table in front of the user.



A single user may use the Holobench interface with head-tracked stereo display, or a number of users may work with head-tracking disabled. The default interaction mode with the *Manufaktur* at the Holobench is to hold the virtual ray-casting tool in the preferred hand while holding the fly tool in the non-preferred hand. Object selection is accomplished by clicking with the ray tool on the object intersected by the light ray, whereas objects are moved by holding down the ray tool button and

dragging the object. While an object is thus grabbed one may hold down the button on the other stylus to rotate and tilt the object using both hands at once. The ray tool uses the so-called go-go technique [20], where the distance of the virtual hand from the user scales non-linearly with the distance to the physical hand. This allows the user to push objects far away and drag far objects near in a single movement. Without the go-go acceleration objects could only be dragged/pushed one physical arms length at a time.

Flying through a scene is accomplished by holding down the button on the fly tool and moving it in the desired direction. Tilting the stylus left or right turns the camera around the vertical axis. If the stylus is turned to a horizontal position (as if balancing a painters palette in the hand), the world-in-miniature view takes effect, and the user is presented with a miniature version of the workspace in his virtual hand. Objects in the miniature world may be selected, moved and turned with the ray tool just like the objects in the normal view. This includes the user's camera that may be moved to change the user's view position from the outside [25].

Pop up menus on objects as well as a global, background, menu are activated by clicking a second button on the stylus in the preferred hand. The menus are presented in 3D space as a rotating wheel from which items can be selected. This menu interface is not ideal, and other alternatives will be explored in the future.

Double clicking on a document container with the ray tool may send a collaboration message to a designated Windows machine that will launch the accompanying Windows application there. Selecting local files for display, changing colours, and entering titles is accomplished through traditional 2D GUI widgets implemented via the Qt toolkit [29].

The Holobench interface is intended as a first step towards a virtual project wall or virtual project room taking advantage of high bandwidth 3D interaction modes to achieve a truly augmented reality interface.

TOWARD ROOMWARE SUPPORT FOR WORKSPACES

In this section, we discuss how we are developing the collaborative *Manufaktur* prototype into a kind of roomware [27] as a work environment which can satisfy the four levels of use identified so far: a spatial environment, parallel environments, collaboration, and bridging between the digital and physical worlds.

The design of the room, which supports projecting larger displays of *Manufaktur* into the physical environment, has three objectives:

- Supplying a resource for peripheral awareness of people's activities and their working materials
- Supporting several collocated people in interacting with the same digital and physical material
- The visual and perceptual integration of digital and physical objects

Awareness and room architecture

(Landscape) architects and industrial designers do not create static objects, but inhabitable spaces or objects for daily use. This implies that the activities of the people who will engage with their work must be anticipated and interwoven with the design. Besides this inherently social aspect of design work, there are other concerns that further intertwine the natural, the material, and the social. The existing material, historical, and cultural context, as well as the laws of nature governing, for example, the physical, chemical, and ecological properties of elements in the design have to be taken into account. At the same time, a conceptual or aesthetic leap is sought.

This plethora of sometimes diverging demands finds expression in the physical space of design studios in the constant manipulation of contextual arrangements of documents and objects. We envisage a scenario where the *Manufaktur* is projected into the physical work environment to allow collocated collaboration to colonize the electronic medium further. However, the nature of the work also suffuses design studios in a broader sense in that their appearance is often aesthetically and spatially tailored to current concerns. Sketches and images on the walls, models on tables, samples on floors and shelves, as well as the flurries of documents and objects assembled in multiple workspaces divide the studio into zones where different sets of concerns and aesthetic intents are present. They also create a relevant mood or atmosphere [4].

A virtual office or project room needs to be sensitive to the flexible and dynamic approach creative professions take when it comes to designing and re-designing their workspaces to accommodate shifting formations of work-teams and materials. The applied artefacts of an interactive room should not only be adaptable to the type of activities going on in terms of functionality, but should also enhance the experience of the users in terms of aesthetics, as one of being in a certain mode/way of working. There is a significant difference between entering the office of an accountant, and entering the office of an architect. Working with atmospheres is difficult, but it is likely to be achieved through sensitivity to the characteristics of the nature of the work and its materials. Augmented Reality addresses this demand. It respects the native environment by being inserted into it without displacing real objects and materials. However, the aesthetic challenge in applying it to the user setting is to achieve as much coherence as possible between the physical environment and its electronic companion.

Interaction techniques

Bringing 3D systems like *Manufaktur* into the physical room raises many design issues regarding interactions with digital objects projected into the environment. These issues have been explored in a series of experiments using various interaction techniques. The interaction techniques span experiments with Spaceballs, gesture recognition, two handed manipulation of VR objects on a Holobench, development of a grabbing/pinching device for augmented reality, i.e. a handheld tool similar to a pair of tweezers [11].

The aim of the experiments is to find a comprehensive and intuitive way of interacting with digital objects in augmented reality. Interaction devices need to be easy and straightforward to grab and use. Having to don a helmet and gloves before one can interact with digital documents and design models impede the fluent transitions required in design work. We will explore further interaction techniques based on unobtrusive and simple devices which are easy to pick from tables and walls, e.g. ‘tweezers’ that aim to emulate or transpose ways of interacting with ‘real’ documents and objects.

CHALLENGES FOR FUTURE WORK

Having developed the software infrastructure and means to provide room oriented display and interaction still leaves us with a number of important challenges to address. We will briefly discuss two of these.

Supporting portable workspaces

As mentioned earlier, the work in architecture and design is characterized by a high degree of flexibility in forming ad hoc workgroups. Work often gets done informally, e.g. in the hallway, dropping by a colleague’s desk, in discussions during lunchtime etc.

These shifts between working modes (e.g. in groups and individually) suggest that being tied to a computer monitor for the display of *Manufaktur* is a cumbersome solution. To make *Manufaktur* a portable virtual workspace manager, which can be activated at different locations in relevant buildings, we need to support portability of documents, models and workspaces as suggested by [21, 22, 27]. When colleagues meet in a *Manufaktur* augmented room in an arbitrary location, they should be able to bring up relevant workspaces on walls or tables without having to go to a keyboard and type URLs or the like. One possible way of supporting this could be to have small databases of workspace IDs stored on palm computers and using IR communication to invoke a *Manufaktur* daemon listening to an IR port in the room. The distribution architecture should facilitate the manipulation of personal and shared workspaces remote from their office or home project room.

Having global access to personal and shared workspaces does not only provide support for casual meetings and planning activities by allowing new tasks, documents (e.g. a sketch) and objects to be related directly to the context of the workspaces in question. It also provides new ways for discussing and assigning tasks during meetings in collocated workgroups. Having brought along their personal workspaces, similar to the way people today bring along calendars, one would be able to insert new tasks, arrange a collaborative workspace layout, drag, drop, and shuffle new documents and objects into dedicated workspaces.

Physical objects and virtual workspaces

The field-studies show that architects and designers use a heterogeneous array of physical materials in their work. To address a diversity of aesthetic, practical, and technical demands, designers draw inspiration and information from

a wide variety of sources. Bridging between digital and physical workspaces calls for physical documents or objects to be linked to materials in one or several *Manufaktur* workspaces, documenting its relations to different contexts.

Currently the physical materials are poorly (or not at all) linked to project documents being worked on in the computing environment. Integration requires both support for computer recognition and tracking of physical material, as well as possibilities for augmenting the material, e.g. by projecting relevant aspects of the computer-based workspace onto the material, in line with the work of [17] augmenting usual paper.

Recognizing physical material in a computer system can be supported at many “bandwidths”: from paper labelling and typing meta data into the system over light weight chips sending an ID over a wireless channel, to full tracking, where the physical object can be recognized at the right physical position relative to virtual objects projected in stereo in the room. Technically all of these recognition methods are possible, but we need to explore to what degree they are practical and economic in a real work environment, e.g. for architects and designers.

As we have seen, enhancing integration between digital and physical objects in an augmented reality environment also calls for a gesture-based interface. The benefits of that would be the emergence of overlapping actions common in the handling of both physical and digital objects. Though the ergonomic and tactile properties of the digital and physical objects are different, the human physical action for moving an object whether it is physical or digital would on a semantic level be very much alike. The idea is that coherence on both a visual and ergonomic level will diminish the notion of working with objects from two different “worlds”. Gesture-based interfaces involving both physical objects and virtual objects raise many new empirical research issues.

CONCLUSIONS

Our approach to CVE is to support work processes. We seek to facilitate users’ collective self-organisation of environments, spaces and traces, which embody the trajectory of their work. They are able to suspend and resume activities, create and recover contexts, and share and transfer their workspaces.

The *Manufaktur* is a collaborative virtual environment, but it is unique in that its focus is on richly representing the disposition of working materials rather than the disposition of people, in terms of avatars.

Finally, we have taken steps to develop *Manufaktur* into a roomware component supporting: peripheral awareness of people’s activities and their working materials; several collocated people interacting with the same digital and physical material; visual and perceptual integration of digital and physical objects

We hope to seamlessly integrate such roomware components in the usual offices and project room environments.

ACKNOWLEDGEMENT

We wish to thank Anders Brodersen and Ina Wagner for their contributions to the work being reported on in this paper. The work has been supported by the EU ESPRIT LTR Project 31870 'DESARTE', and the Danish Research Council's Center for Multimedia (project no 9600869).

REFERENCES

1. Benford, S., *et al.* Broadcasting on-line social interaction as inhabited television. in *Sixth European Conference on Computer Supported Cooperative Work (ECSCW'99)*. 1999. Copenhagen, Denmark: Kluwer Academic Press.
2. Büscher, M., S. Gill, P. Mogensen, and D. Shapiro, *Landscapes of Practice*. Accepted for Computer Supported Cooperative Work: The Journal of Collaborative Computing, Forthcoming.
3. Büscher, M., P. Mogensen, D. Shapiro, and I. Wagner. The *Manufaktur* : Supporting Work Practice in (Landscape) Architecture. in *The Sixth European Conference on Computer Supported Cooperative Work*. 1999. Copenhagen, Denmark: Kluwer Academic Press.
4. Böhme, G., *Atmosphäre: Essays zur ein neuen Ästhetik*. 1995, Suhrkamp Verlag: Frankfurt a. M.
5. Carlsson, C. and O. Hagsand, *DIVE -- A Platform for Multi-User Virtual Environments*. Computers and Graphics, 1993. **17**(6).
6. Churchill, E. and D. Snowdon, *Collaborative Virtual Environments: an introductory review of issues and systems*. Virtual Reality: Research, Development and Application, 1998.
7. Greenhalgh, C., *Spatial Scope and Multicast in Large Virtual Environments*. 1996(NOTTCS-TR-96-7).
8. Greenhalgh, C. and S. Benford, *MASSIVE: A Virtual Reality System for Tele-conferencing*. ACM Transactions on Computer Human Interfaces (TOCHI), 1995. **2**(3): p. pp. 239-261.
9. Grønbæk, K., L. Sloth, and P. Ørbæk, *Webwise: Browser and Proxy support for open hypermedia structuring mechanisms on the WWW*. Computer Networks - The International Journal of Computer and Telecommunications Networking, 1999. **31**: p. 1331-1345.
10. Grønbæk, K. and R.H. Trigg, *From Web to Workplace: Designing Open Hypermedia Systems*. 1999, MIT Press: Boston Massachusetts.
11. Gundersen, K.K. Master Thesis. 1999. Aarhus School of Architecture.
12. Hagsand, O., *Interactive MultiUser VEs in the DIVE System*. IEEE Multimedia Magazine, 1996. **3**(1).
13. Haake, J.M. and B. Wilson. Supporting Collaborative Writing of Hyperdocuments in SEPIA. in *ACM Conference on Computer-Supported Cooperative Work (CSCW'92)*. 1992. Toronto, Canada: New York: ACM Press.
14. Jirotko, M., N. Gilbert, and P. Luff, *On the Social Organization of Organizations*. Computer Supported Cooperative Work, 1992. **1**(1): p. 95-118.
15. Lainer, R. and I. Wagner, Connecting Qualities of Social Use with Spatial Qualities, in *Cooperative Buildings - Integrating Information, Organization, and Architecture*, N.e.a. Streitz, Editor. 1998, Springer: Heidelberg. p. 191-203.
16. Lave, J. and E. Wenger, *Situated Learning: Legitimate Peripheral Participation*. 1991, Cambridge University Press: Cambridge.
17. Mackay, W.E. and A.-L. Fayard. Designing Interactive Paper: Lessons from three Augmented Reality Projects. in *IWAR'98, International Workshop on Augmented Reality*. 1999. Natick, MA: A K Peters, Ltd.
18. Mogensen, P. and K. Grønbæk. Hypermedia in the Virtual Project Room - Toward Open 3D Spatial Hypermedia. in *Accepted for Hypertext 2000*. Forthcoming.
19. OHSWG, Open Hypermedia Systems Working Group home page. 1997. <http://www.ohswg.org>.
20. Poupyrev, I., M. Billinghurst, S. Weghorst, and T. Ichikawa. Go-Go Interaction Technique: Non-Linear Mapping for Direct Manipulation in VR. in *ACM UIST'96*. 1996.
21. Rekimoto, J. Pick-and-Drop: Manipulation Technique for Multiple Computer Environments. in *UIST'97*. 1997.
22. Rekimoto, J. and M. Saitoh. Augmented surfaces: A spatial Continuous Work Space for Hybrid Computing Environments. in *CHI '99*. 1999.
23. Robinson, M., *Double-Level Languages and Cooperative Working*. AI & Society, 1991. **5**: p. 34-60.
24. Schmidt, K. and L. Bannon, *Taking CSCW Seriously: Supporting Articulation Work*. CSCW Journal, 1992. **1**(1): p. 7-40.
25. Stoakley, R., M. Conway, and R. Pausch. Virtual Reality on a WIM: Interactive Worlds in Miniature. in *CHI'95*. 1995.
26. Streitz, N., *et al.* SEPIA a Cooperative Hypermedia Authoring Environment. in *European Conference on Hypertext (ECHT '92)*. 1992. Milano, Italy: ACM.
27. Streitz, N.A., J. Geißler, and T. Homer. Roomware for Cooperative Buildings: Integrated Design of Architectural Spaces and Information Spaces. in *CoBuild '98, Cooperative Buildings - Integrating Information, Organization, and Architecture*. 1998. Darmstadt, Germany.: Springer: Heidelberg.
28. TAN Projektionstechnologie. The TAN Holobench. . URL: <http://www.tan.de>.
29. Troll Tech AS. Qt Product Information. 1998. <http://www.troll.no/qtinfo.html>.