

The Aachen Media Space: Design Patterns for Augmented Work Environments

Jan Borchers

RWTH Aachen University, Aachen, Germany, borchers@cs.rwth-aachen.de

Abstract. *Design Patterns* are a format to capture the knowledge about successful solutions to recurring design problems in a uniform, interconnected, and easily understood way. The format originated in urban architecture, but has made its way into software engineering and Human-Computer Interaction (HCI). This makes them an ideal format to provide guidelines for the design of Augmented Environments, which requires a highly interdisciplinary team to collaborate.

The author, who published the original book on HCI design patterns, has worked on Augmented Environments since 1995. He has captured his experience in designing such environments within the patterns presented here. The chapter first introduces the design pattern format, gives practical tips for reading and writing design patterns, and presents the Augmented Environments from which the patterns in this chapter were distilled. The main part of this chapter then consists of the collection of design patterns for Augmented Environments. These patterns provide practical guidance on how to create an Augmented Environment, with a focus on architectural, interior design, software infrastructure, and user interface considerations. The chapter includes nineteen such patterns connected into a network, as well as suggestions for additional patterns.

1 AEs: A Highly Interdisciplinary Endeavor

Designing user-friendly augmented environments requires expert input from many disciplines: architects for the overall space design, furniture designers for the physical “stuff” in the room, hardware integrators for the interactive devices, software engineers for the OS, middleware, and applications, HCI experts for designing the users’ experience with the interactive technology, psychologists for judging the cognitive impact of the space, and social scientists to support specific activities and work patterns in the environment. Other interactive products usually only require a subset of these disciplines, but Augmented Environments are unique in that they demand the joined expertise of physical and virtual environ-

ments, of physical space design and interaction design, to arrive at successful solutions.

2 Why Design Patterns?

This chapter will review and summarize design recommendations for such spaces, pulled from all of the above disciplines. It will use a particular format for these recommendations, called *Design Patterns*. This is no coincidence: this format is, one might argue, almost destined to be used in this case. The reason is that design patterns, just like Augmented Environments, have developed at the crossroads of traditional architecture and HCI.

The idea of pattern languages was originally conceived by architect Christopher Alexander. Each pattern of his large collection [Alexander, Ishikawa, Silverstein, Jacobson, Fiksdahl-King and Angel 1977] captured a successful solution to a recurring problem in urban architecture, from ideal sizes for working neighborhoods, to the right layout of a street cafe front, down to details such as the distribution of seating areas in a single room. Each pattern carried a descriptive name and described its problem, solution, examples and connections to other patterns on several pages of text and images in a uniform, highly structured format, making it easy for the reader to pick up these design recommendations. The connections to other patterns formed a graph called a pattern *language*. The theory behind pattern languages was laid out in [Alexander 1979], and a sample project using them in [Alexander, Silverstein, Angel, Ishikawa and Abrams 1988].

Alexander actually wrote his patterns not for professional architects, but for the inhabitants, or users(!), of a building or neighborhood, to create an interdisciplinary *vocabulary* that would allow them to express their interests and preferences in the design process in a user-centered way. This idea implements well-known results from psychological research about *verbal recoding*: “When there is a story or an argument or an idea that we want to remember [...], we make a verbal description of the event and then remember our verbalization” [Miller 1956]. The idea can be recalled when its short name is remembered.

Around 1987 [Beck and Cunningham 1987], software engineering picked up and quickly adopted the pattern format, though not quite the larger ideas behind it. The *Gang of Four* book on patterns in software engineering [Gamma, Helm, Johnson and Vlissides 1995] became one of the most widely known programming titles, but unlike Alexander’s, its patterns were not readable by non-experts or users, and they did not adopt Alexander’s inherent goal of creating more human-friendly environments, as he pointed out in [Alexander 1996]. Nevertheless, annual Pattern Languages of Programming (PLoP) conferences have continued to refine the otherwise successful concept of patterns in software engineering.

The “true heir” to the design patterns idea is arguably HCI, not software engineering. Both architecture (in the physical world of buildings) and HCI (in the virtual world of user interfaces) deal with the task of making these environments

as humane to use as possible. Software engineering, on the other hand, is concerned with the construction of robust underpinnings and structures for this user experience—much in the way structural engineering complements the work of the architect.

Surprisingly, Alexander's ideas were referenced as early as 1986 by several seminal HCI books [Norman and Draper 1986], [Norman 1988, p. 229], [Apple 1992]. Still, it wasn't until the late 90s that the concept was explored more rigorously in HCI [Barfield, Burgsteden, Lanfermeijer, Mulder, Ossewold, Rijken and Wegner 1994], [Bayle, Bellamy, Casaday, Erickson, Fincher, Grinter, Gross, Lehder, Marmolin, Moore, Potts, Skousen and Thomas 1998], [Granlund and Lafrenière 1999], [Borchers 2000b], [Borchers 2000c], [Borchers, Griffiths, Pemberton and Stork 2000]. In 2000, the author published the first book on HCI design patterns [Borchers 2000]. Since then, several successful titles have followed [van Duyne, Landay and Hong 2006], [Tidwell 2005].

So far, the most thorough review of Alexander's work and its implications for HCI can be found in [Borchers 2000]. [Dearden and Finlay 2006] compares recent HCI design pattern collections.

3 The Pattern Format

The design patterns in this chapter are arranged such that they deal with design decisions in the rough order that they are faced when planning an Augmented Environment.

Each pattern is presented as several pages of illustrated text, using a very uniform structure and layout with the following components, closely following Alexander's original form [Alexander 1977, p. x]. To better see how this structure maps onto the actual patterns, each component includes the corresponding text from the first pattern in the language.

Name: A meaningful, concise name identifies the pattern and serves to build a vocabulary for design discussion. The first pattern, for example, is called 500 square feet. Following Alexander's example [Alexander 1979] like most pattern books and collections do, pattern names are always written in lowercase small caps in this language.

Ranking: This indicates the validity of the pattern, i.e., in how far the author believes this pattern is an exclusive option to solve the problem described. One star (*) means that it's probably only one of many options; two stars (**) suggest that this is likely among the best of only few options to address the problem, while three stars (***) indicate that any problem of the described kind will likely include an example of this pattern in its solution. The first pattern, for example, is ranked with two stars (**).

Picture: This gives a "sensitizing" and easily understood example of the pattern applied, to quickly help newcomers to get a feel for what the pattern is about, and to set the mood for the remainder of the pattern. The first pattern has an

image of the Aachen Media Space, for example, to illustrate the space requirements that pattern talks about.

Context: This is a single paragraph that explains which earlier patterns the present pattern helps to implement. It gives the designer an idea when, at which stage of the design process, to consider this pattern. The context references back to “higher-level” patterns in the language—in our case this means patterns that the designer should have looked at before the current one.

Together with the references, the context constitutes the edges of the pattern language graph, while the patterns themselves are the nodes of the graph. Generally, context and references are forward and backward versions of the same link, although exceptions to this rule are sometimes made if it appears didactically useful; after all, the language is designed for humans, not computers, so satisfying precise graph conditions does not constitute a value by itself in this case.

In our example, the pattern context says that “you are just beginning to plan your Augmented Environment (AE), wondering where in your building to locate it and how big it should be.”

After the context, the “introductory” part of the pattern is complete, and the “core” part of the pattern begins. This is indicated by three diamonds (◆◆◆).

Problem: A short problem statement starts this core part, by summarizing which problem this pattern addresses. Since design is always a tradeoff, the statement is often worded in the form of competing “forces”, or design tradeoffs. In our example, that statement is the following: **“The bigger the AE, the more flexibly it can be used; but space in most institutions is severely limited.”** The problem statement is always printed in **boldface** because together with the solution (also in boldface), it is the core message of each pattern. If you want to quickly skim a pattern language, you can usually just read the problem and solution statements, and you should have a good idea of what the patterns are about (if the language was written well).

Examples: This is usually the longest section in a pattern. It expands upon the problem statement, and provides examples of existing solutions to it, or other empirical or theoretical background information. It is the “rationale” explaining why the solution provided next is sound. It also often introduces additional recommendations that usually only come “on second thought”—or after the experience of having tried to solve this problem in one or more concrete projects.

Solution: The solution is the single most important pattern component, generalizing the examples into a clear, but generic set of instructions that can be applied in varying situations. If you only read one sentence of each pattern, read this one. It is also printed in **boldface**. Our example pattern suggests: **Reserve a space of at least around 500 square feet for your Augmented Environment. That size will support typical meeting activities for groups of up to fifteen people.**

Diagram: While the initial photo gave a concrete example to set the scene, and was the visual equivalent of the “Examples” section, the *diagram* describes the solution and its constituents graphically in a simple, memorable sketch; abstracting from the details, it is the visual equivalent of the solution statement. When going back through a pattern language you read earlier, the diagrams, due to their visual, iconic nature, should help you to very quickly recognize each pattern again while browsing—faster than by reading any text. Our sample pattern has a simple rectangular floor plan sketch with the words “500 square feet” written into it, reminiscent of what you would find in the floor plan of a building.

This completes the core part of each pattern, as indicated by three more diamonds (◆◆◆).

References: Finally, the *references* point you to subsequent patterns that can be used to implement or flesh out the present pattern further. They are not to be confused with scientific references (citations) to other works. They are like annotated forward links to smaller-scale patterns that come later in the design process; the “context” links are the corresponding backlinks. You can check this: generally, all patterns that 500 SQUARE FEET (1) refers to in its references, and only those, will have a link back to this pattern in their context. Our sample pattern has a fairly long reference section:

“Once the required size is determined, you should consider the right location for your AE—WINDOW SPACE (2). The space that the room takes up can be used much more efficiently if its furniture supports multiple usage scenarios—FLEXIBLE FURNITURE (4) and by making mounts for unforeseen technology to be added later available under the ceiling—CEILING GRID (5) and floor—RAISED FLOOR (6). The size of your room also determines how big your group displays have to be to be readable from anywhere in the room—MEDIA BOARDS (7). If you plan to include noisy equipment or rear-projected wall displays, consider adding a shell around the room for that technology—ROOM SHELL (12).”

4 Pattern Reading and Writing Tips

Much of the advice in these patterns will sound quite trivial to an expert from the respective traditional profession—“of course people want daylight in their environment,” any architect will say, for example, about the pattern WINDOW SPACE (2). The point of this collection is to make people from outside each respective profession aware of these rules as well, in the interdisciplinary way that patterns were originally intended to be used. AEs are a new field that requires interdisciplinary exchange across traditional professional boundaries. Each of us has a lot of basics to learn in this process.

With AEs being a fairly young field, many patterns are still really only *proto-patterns*, with only one or a few examples drawn from the working experience of

the author, as well as other members of the RUFAE network, who authored other chapters in this book.

When using the patterns in the language in order to apply them to your own design of an augmented environment for research or deployment, keep in mind the above tips about skimming problems and solution, and finding patterns again by diagram. It is also important to understand that patterns are not cookbook recipes: Usually, you will still have to make a creative design decision; the patterns just help you to think about the right questions at the right time, and to avoid overlooking important design implications.

After applying these patterns, or maybe based on your existing experience in the field, you may feel that the patterns need amendments, corrections, additional examples, etc. This is the best way to use this language; let it evolve, personalize it with your own examples from projects you and your colleagues know, change or drop some patterns, and add others.

The author has guided students in writing their own patterns in a semester-long class dedicated to HCI design patterns for several years [Borchers 2002], and while reading patterns is very easy (that is what they were designed to support), writing patterns is much more difficult to learn to do well. Some tips to keep in mind:

- Make sure to get the abstraction level right. General “golden rules” such as PROVIDE FEEDBACK are very useful to analyze existing designs and explain problems, but as design guidance they are too generic. A good test is whether the pattern provides some *constructive* guidance as to what to do in the specific context. On the other hand, if patterns become too concrete (AVOID COMBOBOXES UNDER WINDOWS XP, it becomes harder to apply to different interfaces, and becomes outdated as quickly as the toolkit it refers to. One test that may help is to see if the pattern can apply to interfaces of different modalities—not just GUIs, but also voice or gesture interfaces.
- Get the granularity right. Often an initial pattern idea may draw in more and more additional advice about the topic until it becomes too large for a single pattern. A good limit is 2-4 book pages. Beyond that, consider drawing out different aspects of the solution you are proposing, and turning them into a connected, hierarchical set of individual patterns.
- Pay attention to the problem statement. The problem statement often seems harder to write than the solution statement! That is because the problem statement forces you to really think about when your solution is applicable—and when it is not. If you have a hard time coming up with a good problem statement, you may have a one-size-fits-all solution that is not refined enough yet. Similar caveats apply to the context, it also helps you to consider when, or even if, the pattern is useful.
- HCI Design Patterns aim to capture tried solutions to recurring design problems in interactive systems. As such, they have been shown to help especially newcomers to quickly get up to speed in a new interaction design area. They

are less effective to improve the work of seasoned professionals, because those experts already know most of what the pattern language suggests. Use this language, or your own, especially to communicate with customers, others outside your own discipline, and to get new team members up to speed—much more quickly than by having them look at all the projects your company did over the last 20 years.

5 About The Group

To understand what kind of experience the following patterns represent, it is important to know from which background they're coming—what projects have led to this list of patterns, and what were the environments that have been constructed to serve as testbed for the validity of these guidelines.

The author has worked on technology-augmented collaborative work environments since 1995, when he contributed to the conceptual design and setup of the *Conference/Classroom of the Future* [Mühlhäuser, Borchers, Falkowski and Manske 1996] in the *Ars Electronica Center* in Linz, Austria. Subsequently, as acting assistant professor at Stanford University in Terry Winograd's Interactivity Lab, he worked on the *Stanford Interactive Workspaces Project* and its *iRoom* environment, in particular contributing to some of its cross-platform middleware applications, and initiating the *iStuff* project to support rapid physical prototyping of post-desktop interaction techniques. He was also involved in the *iSpaces* project that coordinated and exchanged interactive workspace technologies between Stanford University and KTH Stockholm in Sweden, where the *iLounge* was constructed, and that moved *iRoom* technology into experimental classrooms in Stanford's Wallenberg Hall building.

When joining the CS faculty at RWTH Aachen University to establish a new chair in Media Computing and HCI in 2003, the experience from these projects led to his design of the *Aachen Media Space* that aimed to improve on many aspects of the previous spaces, in particular their flexibility, appropriateness for long-term use, and ability to seamlessly handle audiovisual media in addition to more traditional document types. The design of this space was heavily influenced by Jan's continued involvement with the *RUF AE* research network, in particular through Saadi Lahlou's and Volker Hartkopf's work on highly flexible workspace designs.

The present pattern language therefore reflects the accumulated experience from these projects, and takes most of its concrete examples from the Aachen Media Space as the latest generation of these environments that the author has created, and is continuing to develop, with his team. However, the Aachen Media Space is far from finished. It is a work in progress, but it serves three purposes:

- as a melting pot and testbed to install our own prototypes and new systems to support post-desktop group work on multiple media;

- as an everyday meeting room our group actually uses on a daily basis;
- and as a demonstration and presentation environment to visitors interested in our work on augmented environments.

Thus, using it as an example for the patterns in this language is simply based on observations what has worked well so far.

The language focuses on the patterns the author considers most central. Obviously, there were many more candidates for patterns that did not make it into the this language because of space considerations. Some candidates were:

Mobile Walls — Larger environments can benefit greatly in their flexibility from walls that can be moved around by a single person (but consider their weaker noise blocking qualities).

Tackable Surfaces — Hanging up good old paper is still necessary, even in a hightech environment. Some wall materials will suffer when even the tiniest scrap of sticky tape is applied; avoid them.

Tiling Tables — Tables with wheels really begin to shine if their surfaces can be folded upright to stow them against a wall with little wasted space. Check their robustness though.

Autonomous Seats — Some augmented environments, like the Spaces in Wallenberg Hall, feature chairs that have built-in power supplies and/or networking access. Makes them less mobile but more convenient.

Wired and Wireless Access — Naturally, any AE will offer WiFi access. But wired connectivity has its place, especially when working with large digital media files, streaming video presentations, and similar tasks.

Presentation Wall — Sometimes there's the need for something larger than a mobile Media Board, but not as fancy (and immobile) as an interactive wall. A projection screen mounted on a stand or on a wall, with an ultra-short-throw projector on a table in front of it, will usually do fine. This only concerns larger spaces where people cannot read the Media Boards across the room anymore.

Printing & Scanning — Scanning in an analog document that someone brought to the meeting is often a convenient feature if it is fast, fluid and seamless. A fast feeder scanner, or a mounted digital still camera, serving all files up for the AE to access, is a great addition. Printing documents, on the other hand, is usually better done on a networked printer in an adjacent room to avoid the noise and disturbance.

Cross-Platform Support — The AE needs to handle the machines of regulars and visitors well independent of their operating system. Ensure access to the basic features such as file sharing for anyone coming to the environment with a minimum of hassle.

Recovery-Oriented Computing — The middleware used in the AE should support graceful degradation and self-healing when individual components—even a central server—fail to be reachable. Stanford's iROS, and its commercial successor TeamSpot, are examples of such infrastructures.

Legacy Software Support — Even an advanced AE will face the challenge of visitors “just wanting to show their PowerPoint slides” in everyday use. So make sure your system, as advanced as it may be, does not sacrifice compatibility to widely used software. Not only will this save you from problems with visitors, but also, by supporting standard software and adding to its features through scripting and plugin architectures, you will benefit from their continued evolution and development, instead of having to keep up with those trends through upgrades of its own.

15 People Max. — This is a recommendation on the maximum group size using your AE. Not only will larger groups present problems in managing them effectively, but they will also require a larger space, and in turn larger displays to be readable from a distance, and the social protocols (see the pattern) tend to break down due to anonymity.

Full-Time Technician — Not a design pattern, more a process recommendation. Experience from several different rooms, at different institutions, even countries, and over quite a few years, has confirmed that such a high-tech environment will quickly fall apart if there is not at least the equivalent of one full-time technician who feels that an always-ready, always-working AE and related technology (videoconferencing unit, Media Boards, etc.) is their pride and responsibility (possibly along with maintaining the remaining infrastructure of the group owning the space). Keep this in mind when budgeting.

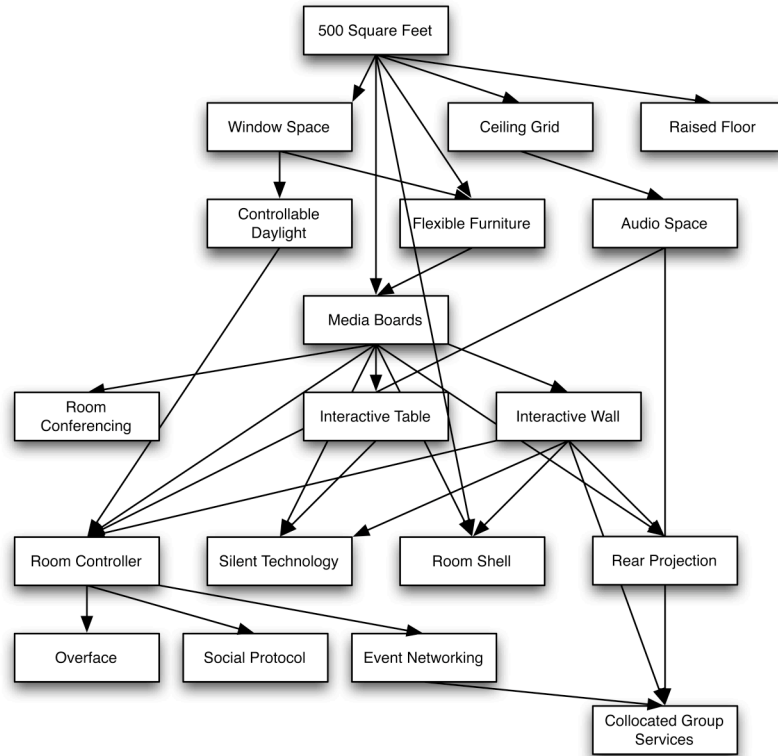


Fig. 1: The Media Space design pattern language.

6 The Pattern Language

1 500 SQUARE FEET **



Fig. 2: The Aachen Media Space in use.

...you are just beginning to plan your Augmented Environment (AE), wondering where in your building to locate it and how big it should be.



The bigger the AE, the more flexibly it can be used; but space in most institutions is severely limited.

Sometimes an entire floor, department or even building is being redesigned from scratch, and there is enough commitment and funding to turn the entire space into an AE (see Hartkopf's chapter on floor designs at CMU). But in most cases, the augmentation is limited to a common space such as a meeting room or semi-public area that all members of a team can use. The deciding factor for its size is the type of activities it should support. A typical meeting-room-sized space is the most common format; it supports one collaborative activity at the same time. To support multiple collaborative activities at the same time will require a much larger space to avoid distracting acoustic interference between the groups.

This also means that such a room can typically support teams of around five full-time members, with roughly another ten infrequent visitors who come to some of the meetings. The tasks envisioned for spaces of this size include internal

presentations, brainstorming and group discussions, administrative meetings, and hosting external visitors and their presentations, and presenting demonstrations to them.

The Stanford iRoom has 20'3" x 17'10" = 360 square feet of usable main room area, plus another 210 square feet occupied by rear-projection displays on two sides, for a total of 570 square feet (including the technology hidden behind walls). It also has an adjacent developer area with workstations covering another 250 square feet. The KTH iLounge has a similar size. The Aachen Media Space is 8 by 6 meters, or 517 square feet. Our latest Media Spaces for the HumTec center at RWTH Aachen University are more varied, and will accommodate activities of smaller and larger groups: 7 by 4.6 meters, 9.4 by 4 meters, and 5.8 by 14 meters, or 350, 400, and 870 square feet.

Therefore:

Reserve a space of at least around 500 square feet for your Augmented Environment. That size will support typical meeting activities for groups of up to fifteen people.



Once the required size is determined, you should consider the right location for your AE—WINDOW SPACE (2). The space that the room takes up can be used much more efficiently if its furniture supports multiple usage scenarios—FLEXIBLE FURNITURE (4) and by making mounts for unforeseen technology to be added later available under the ceiling—CEILING GRID (5) and floor—RAISED FLOOR (6). The size of your room also determines how big your group displays have to be to be readable from anywhere in the room—MEDIA BOARDS (7). If you plan to include noisy equipment or rear-projected wall displays, consider adding a shell around the room for that technology—ROOM SHELL (12). ...

2 WINDOW SPACE ***



Fig. 3: Daylight entering the Aachen Media Space.

...you have determined the size of your AE—500 SQUARE FEET (1), and are wondering where to put it in your building.



An AE is often added to an existing building or department, and the remaining available space is often unattractive with no daylight; but humans do not enjoy being in a space without daylight for long, especially for the creative activities that are supposed to happen in the AE.


When the Stanford iRoom was planned, the only location with enough contiguous space left for it was in the basement, so it was installed there. But when prolonged meetings were held there, participants frequently expressed an urge to leave the space and continue outdoors in daylight. The space seemed fine for solitary, concentrated work in which one became largely oblivious to one's surroundings, but for the social activities for which it was designed, a more humane environment with daylight turned out to be important.

Similar observations were made with users of the Classroom of the Future at the Ars Electronica Center in Linz, Austria (AEC). Short-term visitors had no problems with its artificially lit interior, but people staying for more than an hour would frequently report slight discomfort.

As a result, the location of the Aachen Media Space was chosen such that one entire wall was a window front, allowing lots of daylight to enter. The users of this space have not reported any of the adverse symptoms observed in the other spaces. The EDF space in Paris also features more frequent windows.

Therefore:

Make sure your AE has as much access to daylight as possible, ideally directly through windows, or at least indirectly via indoor glass panes and ceiling lights.



Unlimited daylight can make it hard for people to work with the technology in the space, so make sure to make its intensity adjustable—**CONTROLLABLE DAYLIGHT** (3). Some setups will work best with daylight coming from a different direction than others, so keep your interior setup moveable—**FLEXIBLE FURNITURE** (4). ...

3 CONTROLLABLE DAYLIGHT ***



Fig. 4: Half-closed blinds in the Aachen Media Space.

...you have found a space that provides daylight—WINDOW SPACE (2), but are wondering what kind of blinds may be necessary.



Daylight is essential for humans to feel comfortable in a space for prolonged periods of time. Yet, some displays may seem to require low-light conditions to work best.

When projection technology was just becoming affordable for office use, its light output was often insufficient to work in daylight conditions. This has changed, and most projectors and screens actually work well even with indirect daylight on them. There is hardly any need for darkening a room for projection, unless the projector is underpowered for its projection surface (700 ANSI lumen per square meter is a good rule-of-thumb for projectors).

The Aachen Media Space has a whole window front on one of its sides, and it has light, electrically controlled vertical blinds that can be opened, closed and turned to shield the room from direct sunlight, while still letting in some daylight.


The environment at EDF is located in a room with an extensive glass front. Much of that front is shielded from daylight using heavy shutters to allow projection screens to be read easily. But there are easy ways to open those shutters from within the room, and an additional area high up below the ceiling can be opened to daylight as well.

The Media Space at KTH Library in Stockholm also features large overhead daylight windows, but it is equipped with drapery that can be remotely controlled

from a central room controller to automatically open or close, depending on the usage scenario selected.

Therefore:

Equip your daylight windows with easily adjustable blinds. You should be able to close and open them fully in 10 seconds. They should allow at least three different settings with different daylight levels. When closed, they should leave the room inside light through their own color. Avoid black drapes.



You may want to consider making the blinds electrically controllable and integrating them into the room control interface—ROOM CONTROLLER (15). But make sure that they close and open quickly: if you have to wait for 2 minutes before they have opened fully, nobody will bother using them much. This can be difficult with electrically controlled systems. ...

4 FLEXIBLE FURNITURE **



Fig. 5: Moving a wall in the Augmented Environment at EDF in Paris.

...you have determined the general size and location of your space—500 SQUARE FEET (1), WINDOW SPACE (2), and are beginning to think about how the room should be equipped with furniture.



An AE will be used for a variety of different tasks; however, these tasks all require different furniture configurations to be supported as well as possible.

Activities that the Aachen Media Space, for example, has been used for, include internal and external research talks and presentations, brainstorming sessions, group meetings to discuss progress and plan next steps on a project, discussion rounds with a mix of 5 to 10 student or faculty members, two-way and multipoint video conferences with one or multiple participants at both the local and remote end or ends, and technology demonstrations of current research projects to internal or external audiences. The other AEs built by RUF AE members have mostly seen a similar breadth of activities.

Almost every single one of the above tasks, however, requires a different spatial arrangement of chairs, tables, displays and other technology for optimal support. The RUF AE environments have addressed this need to varying degrees, and in different ways.

The Stanford iRoom contains a large central table that includes rear-projection, but is not moveable. Three fixed SMART Boards are built into one of

its walls, and a fourth, 10-Megapixel display wall into an adjacent wall. The result of this arrangement was that it supported small group discussions quite well, but it was difficult to use the room for presentations, as the table in the middle did not leave enough room for additional people to sit in front of it towards the wall displays. It also created a boundary that people were unlikely to cross in order to engage in interaction with a wall display in collaborative work sessions [Gill and Borchers 2003]. The KTH iLounge followed a similar design, with an immobile large table in the middle of the room, and a display integrated into one wall.

The spaces at EDF, CMU, and Aachen were designed to be more flexible. Their furniture can be moved and even folded away easily. Everything is on wheels and easy to move by one person alone; in Aachen, this extends not only to tables and chairs, but even the wall displays—*MEDIA BOARDS (7)*. EDF even has a larger space with moveable walls, which can be pushed around by one person, yet provide perfect visual and decent sound shielding.

The Aachen Media Space often looks a mess after a meeting, because tables and chair are all over the place. However, they are placed that way because when people came in for the activity, they grabbed a chair *and a one-person table* and wheeled them to where they wanted—which hardly ever leads to a straight, perfectly aligned square or similar formation. And before representative events, one of our staff members goes in and arranges the room to fit that event as well as possible in advance—which is easy because he can move everything alone and without effort.

The one-person rule is critical: Not only should a table have wheels, but one person, male or female, with more or less physical strength, has to be able to move each piece of furniture with one hand, ideally even stow and unfold tables, because the other hand is often holding a laptop or notepad as he or she comes in. Anything that does not fulfill that criterion, from our experience, will not be moved around most of the time even when it would be beneficial to the task at hand.

Tables that can not just be wheeled around but also folded up and stowed away easily will be somewhat less sturdy than their inflexible counterparts. So expect to pay significantly more for the “mobile” version of your furniture, in order to get something that is built as sturdy as you need it—tables to lean or even sit on, for example. Good mechanics make the difference here. A clunky folding mechanism that gets stuck frequently will never get used and therefore will be a waste of your money.


For the Aachen space, the most helpful question to ask in each step of the design and furniture was, “will this nail down any spatial aspect, the location or orientation of any activity in the room? ”. When the answer was “yes”, we looked instead for a more flexible solution. In general, anything that would get fixed to a wall or ceiling or would be too heavy to move easily, was a candidate for improvement. Because of that rule, we have not even installed a fixed projector and

screen in the room, and have instead opted for moveable large screens—MEDIA BOARDS (7).

Similar ideas about flexible furniture were used by the author in an earlier project merging future conference and classrooms [Mühlhäuser et.al. 1996].

Therefore:

Try not to bolt down any aspect of the room's spatial furniture and technology arrangement. Put as much of your furniture in your AE on wheels, or make it otherwise easily moveable by a single person: Chairs, tables, even displays and other equipment. Make tables fold and stow in a corner easily, saving space when not in use.



Putting group displays on wheels as well is a good idea—MEDIA BOARDS (7).

...

5 CEILING GRID **



Fig. 6: The Aachen Media Space ceiling.

...you have decided about the general size of the space—500 SQUARE FEET (1), and are now wondering where to run the cables that the room requires.




AEs require a lot of cables to interconnect the technology in the room. But running cables across the floor is hazardous. Also, new small devices like additional switches, power adapters, microphones, cameras or other sensors are often introduced into an AE after the fact, but there is no easy way to mount them anywhere.

The Aachen Media Space features a 3-foot ceiling grid of open “cable trays” made from strong wire mesh. It has served to hold the initial audio cables for the speakers in the room—AUDIO SPACE (14), as well as additional optical and copper network cables, media converters, switches, power adapters, a wireless router and other parts introduced to the environment as technology progresses. Having this easily accessible structure under the ceiling makes it much more likely that new technologies are installed promptly, without requiring any construction work or drilling holes. It also helps to get technology off the floor, reducing tripping hazards and extending the life span of cables. All the time, this grid usually remains unnoticed by people using the room, making the technology hide from their conscience.

The space at EDF features a similar grid.

Therefore:

Put a wire mesh grid under the ceiling of your AE, at about a 3-ft raster, to have a quick way to attach or hold cables and small devices installed in the room.



The grid is a good place to hold, for example, cables for room speakers—
AUDIO SPACE (14) ...

6 RAISED FLOOR *



Fig. 7: An open floor panel in the Aachen Media Space. The entire panel including the floorbox can be moved.

...you have decided on the overall size of your room—500 SQUARE FEET (1), and are wondering where to put the cabling, switches, and related technology infrastructure.



Users in an AE want access to power and networking where they sit. Floorboxes are a natural choice for this. But in a flexible space, people end up moving their chairs and desks into many different spatial arrangements. Even with wireless networking, users will usually still want access to power for their laptop or other personal devices, especially when presenting or in longer sessions.

An elegant solution to this problem is a raised floor. The EDF lab features a reconfigurable raised floor with tiles that can be taken out completely. Most major cabling is done in the raised floor. The Stanford iRoom uses a similar arrangement, although it is less accessible due to the immobile furniture.

The Aachen Media Space also features a raised floor, and here the tiles are not only removable, but the tiles with floorboxes in them can actually be discon-

nected completely, allowing users to even move the floorboxes around, e.g., to cluster power and fast networking access points in a corner of a room to set up a cluster of computers or neighboring displays. The connections providing power and networking to the floorboxes are numerous, long and flexible enough to be moved elsewhere in the room before being re-connected to the floorbox at the new position.


Floorboxes, even when they are completely flush with the ground as in the Media Space, always constitute a tripping hazard when cables protrude from them. However, with moveable floorboxes this risk can be reduced: The floorboxes will end up spread out across the Media Space floor in an arrangement that best fits its current typical layout. For example, we have moved one floorbox each directly beneath where we typically have our Media Boards, thus moving their power connection out of the way.

The floorboxes ideally contain mains power and Ethernet ports. They can also have rolled-up Ethernet cables inside to save users from having to find those elsewhere. Or, if the users have fairly homogeneous technology (as is the case in the Media Space), power adapters for the standard laptop brand (Apple in our case) can be put right into each floorbox, simplifying the most frequent access needs even further.

Note that you'll have to work with your institution's safety department to ensure that fire safety regulations are being followed (you may need to install separate smoke or heat detectors in the floor), and that electrical safety standards are met if floorboxes are unpluggable.

Therefore:

Raise your floor by about 1ft (30cm), and put power and networking cables/switches into the raised floor. Provide power and network outlets at floorboxes, at least one every 5ft (1.5m) in each direction, and make them moveable to accommodate future spatial configuration needs. Make the floorboxes deep enough to be able to include network cables and power adapters inside the floor boxes.



This is a basic pattern with no further references within this language.

7 MEDIA BOARDS *




Fig. 8: A media board in the Aachen Media Space.

...you have determined the basic furniture to be put into your AE, and decided to make it as flexible as possible—FLEXIBLE FURNITURE (4). You are now thinking about the right display technology to go with this design.



Any computer-supported group activity greatly benefits from large, shared display surfaces that can also be sketched upon and interacted with. However, projectors and projection screens are noisy and hard to move around quickly.

Groups need group displays. We have all huddled over a laptop with others, or squinted at it from across the table at some point to see what someone is trying to show us, and it's not a very productive spatial arrangement.

The most frequent solution to this problem are projectors. They are cheap, can project large images (if they have strong lamps and short-throw lenses), are easy to carry around, and can project on a white wall if necessary. However, projectors have a few drawbacks: First, they invariably include a fan that will create background noise in the room unless shielded from the users. Second, even though


they are quick to move, setting them up and aligning them for good projection takes time, breaking the flow of a meeting. Third, when portable, the projection is usually from the front, and the presenter will often have to look into the lens, and cast a shadow. Projecting from the rear solves that problem, but this only works where there is a special rear-projection screen surface installed, takes up space behind the screen, and provides only narrower viewing angles. Finally, projectors usually require a cooldown time before disconnected from power, further slowing down any quick reconfigurations.

The Aachen Media Space instead features four moveable, interactive forty-inch displays. We called them *Media Boards* because they serve to interact with different media types: They let users project and interact with on-screen presentations, a web browser, other desktop applications, or special whiteboarding software. Each is constructed from an LCD screen with a touch-sensitive SMART Board overlay that lets users use their fingers or passive plastic pens to point and write on the board. The board is mounted on a mobile four-wheel stand that balances the weight of the screen to prevent the board from toppling over when it is being moved. A laptop is placed on a separate shelf under the display, providing a keyboard when needed, and (thanks to its battery) allowing the board to be moved around with power disconnected without having to reboot the computer. A second shelf can be attached to hold, for example, a portable video conferencing unit. The display can be adjusted in its height to serve different purposes, such as presentations (at standing height) and video conferences (at sitting height). Local speakers are mounted on the outer ends above the display. Thanks to wireless networking and a power strip and cable mounting grid at the back of the stand, the only cable going away from the display into the ground is a single power cord. Separate ethernet, DVI, VGA and audio cables are rolled up at the back of the stand to quickly connect other sources such as a visitor's laptop or demo machine when needed.

Plasma displays are less suitable for use in Media Boards than LCD screens, because they are heavier, typically require integrated fans for cooling—SILENT TECHNOLOGY (13), they suffer from burn-in problems with static content (unlike LCDs) and brightness loss over time, and their higher contrast is not crucial for the typical work tasks and lighting conditions in AEs.

Therefore:

Provide mobile LCD screens at least 40 inches in diameter, complete with a hooked-up laptop and optional connections for external sources, with speakers and on a stand that can be wheeled around and is height-adjustable, with a single power cord to plug in for use.



As display technology advances, equipping entire walls and even whole rooms wall-to-wall with interactive display surfaces INTERACTIVE WALL (8) is gradually becoming affordable. But even in such environments, moveable displays are useful, as any area in the room (not just along the walls) can be picked to display and interact with information. The flexibility rule suggests providing display space *as a commodity*, easily made available wherever it is needed at the moment. Note that without a heavy counterweight, it's impossible not to have the stand reach out beyond the front of the screen. One thing to keep in mind, then, when choosing a stand is that people will want to walk up to the display to use it, so its feet should not get into the way. Also, the displays may need to be pushed together as closely as possible for certain multi-screen scenarios. If the screens are to be tilted slightly towards each other when standing beside each other, then make sure that the wheels and sides of the display stand do not get in the way to do so at the required angle. Feet that can be turned towards the inside or outside can help to minimize the impact of this for different scenarios. The forty-inch size of the existing Media Boards in Aachen is sufficient for small group interaction, although it can sometimes be taxing to view smaller details in presentations on them from all across the room. Affordable LCD screen sizes are continuing to increase, however, and sixty-inch screens will be more than sufficient for a typical-size augmented meeting room—500 SQUARE FEET (1).

As with all technology, ensure that your boards don't create distracting noise through fans or transformers—SILENT TECHNOLOGY (13). If you need to consider projectors, choose a rear projection setup—REAR PROJECTION (11). You may be able to address space requirements and remaining noise problems of this arrangement better if you put it behind interior room walls—ROOM SHELL (12). A height-adjustable media board will work well for room-scale videoconferencing—ROOM CONFERENCING (10).

For some tasks, a horizontal display may be better suited, especially in addition to wall displays—INTERACTIVE TABLE (9). The Media Boards should be controllable from a web page and dedicated mobile device in the room—ROOM CONTROLLER (15). ...

8 INTERACTIVE WALL *



Fig. 9: The DynaWall® by Norbert Streitz' group.

...you are deciding on the display technology to put into the AE. You are either planning to turn all walls into display surfaces, or you know that the AE will have only a few, clearly defined, constant tasks. Or maybe you need a resolution higher than what mobile displays can provide. Because of this, you have decided to look into a static installation of display technology instead of, or in addition to, using movable interactive group displays—MEDIA BOARDS (7).



Groups need group displays to work with digital data, and complex information requires lots of screen real estate; however, standard displays come in small sizes and low resolution

Large interactive walls have a big benefit over smaller, mobile displays in that they can display more context and detail at the same time. Walls also are more appropriate than tables for many tasks since they provide a common, easy-to-read vertical perspective on the contents for all users in a collaborative group, whereas tables will need to provide special mechanisms to re-orient content for different users around the table—INTERACTIVE TABLE (9).

The Darmstadt IPSI environment featured one of the earliest interactive walls, the *DynaWall®* [Streitz, Geißler, Holmer, Konomi, Müller-Tomfelde, Reischl, Rexroth, Seitz and Steinmetz 1999], made up from three rear-projected large touch screens built side-by-side into a wall. The Stanford iRoom featured a very similar setup, using three SMART Boards. It was augmented by a fourth, large


10-megapixel display made from 4x3 standard projectors hooked up to a rendering cluster [Guimbretière, Stone and Winograd 2001], with software focusing on brainstorming activities. The three-display iRoom design was subsequently adopted in the design of the Wallenberg Hall interactive learning spaces at Stanford. The iLounge at KTH and the EDF space in Paris also have one wall with a large built-in projection which is usually where presentations are being shown. More recent interactive wall work from the IPSI group includes the *InfoRiver* and *Hello.Wall* [Streitz, Magerkurth, Prante and Röcker 2005] that also focus on informal group communication.

There is not much use for a gigantic screen (beyond around 40 inches) in a meeting room unless its resolution also goes well beyond 1024x768. Some large displays feature high resolution already, another option is to tile smaller standard displays together into a large, smooth surface. Aim for very small seams between tiles for visual continuity in this case. A good target resolution is at least 60 dpi, which is close to the 72 dpi of traditional desktop monitors.

Making such surfaces interactive usually required attaching an external sensor system (such as an eBeam® or Mimio® system) to the display, which then sensed touch via special pens. Alternatively, touchscreens such as SMART Boards can be tiled, with the advantage of supporting bare finger input. More recently, infrared light has been used successfully to accurately detect multiple-finger touch on rear-projected surfaces [Han 2005] and on flat LCD screens [Hodges, Izadi, Butler, Rrustemi and Buxton 2007], opening up the potential for more natural interaction gestures.

Therefore:

Include a wall-size display in your AE. Make it rear-projected or use LCD screens. If budget allows, aim for a pixel resolution of at least 60 dpi, if necessary by tiling multiple smaller displays with as small seams as possible. Make the surface interactive with at least 1mm resolution to allow for touch input and sketching, if possible with multi-touch support.



Middleware can link independent displays together to effectively provide a larger, contiguous display to the user—COLLOCATED GROUP SERVICES (18). This software can help integrate multiple small mounted screens into an interactive wall, or to combine mobile displays into one. Like the Media Boards, the Interactive Wall should be controllable from a web page and dedicated mobile device in the room—ROOM CONTROLLER (15). If projectors must be used, consider projecting from behind the wall—REAR PROJECTION (11), and/or putting the wall technology into a space behind the room's interior wall—ROOM SHELL (12). ...

9 INTERACTIVE TABLE **



Fig. 10: The multitouch table, designed by David Holman from the author's Media Computing Group, in use at a research exhibition.

...you have decided on your vertical display space—MEDIA BOARDS (7), but are considering the need for horizontal display surfaces as well.



Some activities, such as collaborative photo sorting and analysis, are best done by a group around a table. However, traditional laptop or wall displays do not lend themselves well to this situation

Streitz et al. created a touch-sensitive table as part of their *Roomware*[®] research and product development. [Streitz et al. 1999]. The Stanford iRoom included a table, but it was immobile, with its rear-projection built into the raised floor. The KTH iLounge also featured a fixed table in the middle of the room.

More recently, the author's Media Computing Group has built multitouch-sensitive tables that use rear projection and afford more natural interactions with multiple fingers and hands. This technology has been used, for example, in a multitouch table exhibit for a Dutch children's museum. Microsoft's *Surface* project builds on these ideas to create a commercially viable multitouch table.

The big caveat with tables is their inherent orientation. Reading text together is much harder around a table than on a shared wall display. The direct interaction with artifacts that would also be shuffled around on a real table, however (such as photos), is a big plus, especially if the table is multitouch-enabled (single-touch inhibits the natural parallelism and overlap of collaborative human activity). Toolkits such as DiamondSpin [Shen, Vernier, Forlines and Ringel 2004] allow for easier development of rotation-enabled interfaces.

The second issue with tables is their immobility. Heavy technology such as a large, possibly stationary rear projection system should be avoided. Small projectors, or LCD screens are a better choice to keep the table moveable.

Therefore:

Consider adding an interactive group table to your room, especially if you are interested in this topic for research purposes. Choose a design that remains mobile and that is supported through software libraries to develop and run the somewhat idiosyncratic UIs required for it.




Table users are very close to the table, in order to reach its surface. Because of this, it is important that the table does not create unnecessary noise, e.g., through fans in its rear projector or the connected computer—SILENT TECHNOLOGY (13). ...

10 ROOM CONFERENCING **



Figure 11: Users in front of a Polycom videoconferencing unit.

...you are putting together the list of services and technologies to put into your AE—MEDIA BOARDS (7), and are thinking about video conferencing connectivity to the outside world.



Today, video conferencing is used regularly to avoid having to fly people around the globe. However, users often report inadequate audio and video quality, and an overall unsatisfying experience when using simple videoconferencing solutions

It is vital to distinguish between two groups of systems. They both provide a live audio/video link between two sites, but their intended use is fundamentally different:

Personal, or Desktop Videoconferencing Systems: These are typically software solutions that use the processing power, display and speakers of a desktop or laptop computer with an attached camera and microphone (or a headset). They often even support application or document sharing across the video link. Examples include Apple's iChat A/V that comes free with their operating system, or Polycom's PVX® software. They are designed to capture *one person's* audio and video input well. This is their important limitation: As soon as a group of people in one place tries to use these systems, they begin to break down due to lower video resolution, insufficient room characteristics of microphones and limited echo cancellation.

Room Videoconferencing Systems: Their optical camera quality, video resolution (moving to HD as of this writing) and compression technology, microphone design, echo cancellation and audio processing (all done in hardware) are designed to handle multiple speakers in the same room on the system's end. They feature locally and remotely controllable cameras, multiple microphones, and web streaming interfaces for additional passive participants. Prices for high-end systems lie in the \$10k range. A current example is Polycom's VSX® 7000 system.



Getting a desktop system for budget reasons when a room system would have been needed is a certain recipe for failure when trying to use the technology constructively.

For best eye contact, lower the screen showing the remote participants so that their eyes on-screen are at eye level of the local participants (around 4 feet high), and mount the video camera directly above the upper edge of the screen. If the system provides a picture-in-picture feedback view of its own local camera image, enable it to allow participants to notice when they are off-screen, for example. Place that view at the top center of the screen if possible, immediately under the camera.

If frequent collaborators are already known, then the system should be chosen to be compatible with existing technology on the other end. The camera should always be as close as possible to the top of the screen for best mutual eye contact. The Media Board or other screen should be lowered so that participants see eye-to-eye.

Therefore:

Provide a portable but high-end room videoconferencing system of adequate quality for the size of your AE. Put it on top of a MediaBoard screen lowered to sitting eye height, so it can be moved easily.



This is a basic pattern with no further references within this language.

11 REAR PROJECTION **



Fig. 12: SMART Board with rear projection. (Image copyright 2001–2008 SMART Technologies ULC. All rights reserved.)

...you are planning the outfitting of your AE with displays—MEDIA BOARDS (7), INTERACTIVE WALL (8), and are wondering whether to use front- or rear projection for your projector-based setups.



Projectors, even short-throw systems, need a certain distance from the screen. However, in an AE users should be able to interact with displays by touching them, in which case they are directly in front of the display and in the projector's path, casting shadows just where they are working.

Many systems have recognized this problem and opted for rear projection, such as our conducting exhibits [Borchers, Lee, Samminger and Mühlhäuser 2004], even when the user is not directly touching the display. The reason is that rear projection reaches higher display contrast than front projection, improving the visual quality of the interface.

With interactive surfaces, many solutions such as a SMART Technologies' rear-projection SMART Boards, opt for rear projection to avoid shadows on the sur-


face. If there is no space for rear projection, an ultra-short-throw front projection (projection distance to projection width $\ll 1$) can be used.

A more elegant solution may be using an LCD screen, as it does not bring the depth requirements with it (or noise problems) that rear projection has. Therefore, rear projection is usually chosen when relatively low resolution, but a large projection area are needed.

Rear projection does not mean that the projection needs to be mounted in a fixed wall, as the movable SMART Boards demonstrate. In any case, the rear-projection housing may serve to shield some of the projector's noise.

Therefore:

If projectors must be used in an AE, use rear projection for shadow-free operation, better contrast, and somewhat less noise.



This is a basic pattern with no further references within this language.

12 ROOM SHELL

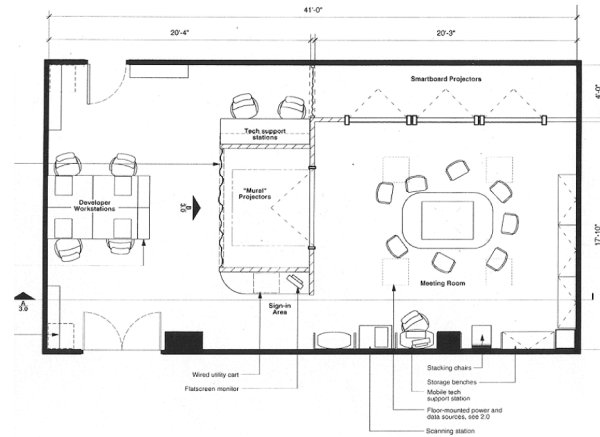


Fig. 13: Layout of the Stanford iRoom. Note the shell spaces for rear-projection technology along the top and left sides of the main room.

...you have decided on the size of your interior space—500 SQUARE FEET (1), and on what display technologies you intend to provide—MEDIA BOARDS (7), INTERACTIVE WALL (8). You are now wondering where to put the extra technology required.



AEs require additional computers, amplifiers, and similar technology that normal workspaces do not feature in the same amount. However, there is usually not enough space to put that technology into the room itself, and it may constitute a noise problem; on the other hand, the technology has to be close to the room because of cable length limitations and the need to access the technology when problems arise.

This conflict was addressed in the Stanford iRoom by putting an entire shell around two sides of the room. That shell was deep enough to house rear projected screens, and even opened up towards a developer space on one side.

The KTH iLounge used a similar approach, putting much of its technology behind a rounded inner “cocoon” that was used by the participants as meeting space.


The raised floor in the Aachen Media Space can be seen as a shell, but its design without projectors and noisy technology removed the need for such a shell, which in turn left more space for the actual room itself.

At UCSD's Distributed Cognition and Human-Computer Interaction Lab, a space adjacent to their meeting environment both houses the rear projection and provides additional storage space for technology.

At EDF France's LDC space in Paris, the raised floor also serves as a shell, providing space for various infrastructure technology. It is also air-conditioned, with adjustable floor outlets, providing the flexibility that moveable walls and furniture require.

Therefore:

If rear-projected screens are needed inside the walls, or noisy technology needs to be stowed away while still being accessible, and your AE does not become too small by this, consider closing off a secondary space, about 4ft deep, around one or more of the walls of your AE.



This is a basic pattern with no further references within this language.

13 SILENT TECHNOLOGY **




Fig. 14: A computer fan.

...you have decided on your overall devices to put into your AE, such as displays—MEDIA BOARDS (7), INTERACTIVE WALL (8), INTERACTIVE TABLE (9)—and are selecting particular models and technologies



Background noise elevates stress levels and is detrimental to users' health, ability to concentrate, and productivity. However, AEs require a lot of technology to be integrated into the space that users work in.

Desktop computers and towers, projectors, plasma displays, external hard drives, and many other devices today contain fans for cooling or otherwise produce noticeable noise levels. Even these low noise levels have been shown to introduce stress to people in the environment.

To avoid these factors, selecting the right technologies is the most helpful step. For example, unlike projectors, LCD screens are silent because they have no fans. Laptops, unlike desktop or tower computers, will run without their fans spinning up for most of the time, and can easily master most tasks in AEs today.


Even within a given technology, there are sometimes huge differences in noise levels between competing products. It pays to check for the dB level rat-

ings. And especially hard drives can develop whining noises with age, so exchange them after two to three years.

The equipment in the Aachen Media Space, for example, was selected to contain not a single constantly running fan, mostly by choosing LCD screens over projectors, and laptops over stationary machines.

Therefore:

Keep the noise level added by the technology to a minimum, by avoiding equipment with constantly running cooling fans, drives, or other noisy mechanics.



If noisy equipment is unavoidable, consider shielding it inside larger encasing, or behind a wall—ROOM SHELL (12). Of course, the original position of the room would have to support a quiet environment to begin with—WINDOW SPACE (2). Also have an eye (or rather ear) on any airconditioning system in the room.

...

14 AUDIO SPACE *



Fig. 15: One of the eight speakers around the Aachen Media Space.

...you have settled on the furniture and display equipment—MEDIA BOARDS (7), INTERACTIVE WALL (8), and are now thinking about the audio side of your setup.




Laptops and stationary machines will often need to play back audio for the participants in the room, but running audio cables across the room is an unsatisfying solution.

An early part of the software development going into the Aachen Media Space was the *AudioSpace* software, originally developed for Mac OS X 10.3 by Stefan Werner in a co-advised diploma thesis [Werner 2005]. It consisted of two parts: You installed a client-side driver on your laptop that would let you choose a new virtual audio output to direct your audio to. This driver would not play back your audio, but instead package it up and send it via a TCP network connection to the *AudioSpace* server application. The server would unpack the data and play it back via its local audio interface, typically connected to the PA system in the room you're in. The result: Anybody in the room could play audio using the room audio system, from their laptops, without any audio cabling to the laptops required.

Other software exists that allows for a similar wireless audio link. The downside, in all cases, is of course that software needs to be installed on each machine that wishes to participate; but the advantage of cable-free access to room audio is quite significant.

Therefore:
Make software available for machines in the AE to wirelessly send their audio output to a room audio system.



This functionality could be wrapped into the general software required to connect to the room infrastructure—COLLOCATED GROUP SERVICES (18), and control could be offered via a mobile device—ROOM CONTROLLER (15). ...

15 ROOM CONTROLLER **




Fig. 16: The UI for a tablet- or PDA-size room controller, showing a room layout with lights, projectors etc., to control the Stanford iRoom.

...you have settled on the interactive technology to be installed in the AE—
 CONTROLLABLE DAYLIGHT (3), MEDIA BOARDS (7), INTERACTIVE WALL (8),
 INTERACTIVE TABLE (9), AUDIO SPACE (14). Now you need to design a way for
 users of the room to control the technology.




Technology added to a work environment often dictates its own, proprietary user interface through its physical design, and is usually only controllable locally at the device, or through a single physical remote control. However, users think in different, more space- and task-oriented ways when trying to operate this technology, and want to operate it wherever they are in the room.

In the Stanford iRoom, three projectors, a high-resolution projection wall, a rear-projected table and multiple ceiling lights had to be operated in order to get the room into a particular state. To achieve this from anywhere in the room, their proprietary control interfaces were adapted to a TCP-based protocol, and a

graphical user interface showing all devices in a simplified floorplan of the room was generated as an interactive web page. Any browser could open that page, and it was adaptable to displays of different form factors (such as mobile devices), but it was also displayed on a dedicated wireless tablet PC that was always available in the room. This allowed users to grab the device or have it passed to them when they wished to change something, with everybody else realizing what was going on.

Therefore:

Make sure your devices are controllable via a standard, open protocol such as serial or TCP/IP communication. Provide a graphical spatial floor map that shows where your interactive devices are situated in the AE, and that lets the user operate the devices from this map. Create the floor map from a service description so that it can be adapted to different display sizes and UI form factors, and so that it can dynamically change as new devices become available. Make the map displayable on any browser, but also provide a dedicated mobile touch display in the room that always shows this map.



The room controller requires the room technology to be accessible via the network—EVENT NETWORKING (17). It can be a good representation for the unified interface to the room and its services—OVERFACE (19). Its usage, as with all other technology in the environment, should not get into the way—SOCIAL PROTOCOL (16). ...

16 SOCIAL PROTOCOL ***



Figure 17: Passing on a mouse for a group display.

...you have picked your hardware to control the room and its services—ROOM CONTROLLER (15), and now need to decide how the technology is operated by the users.



Interactive technology likes to be told when something happens or when it is supposed to do something. But people easily forget that extra step, especially when in the middle of a high-energy brainstorming session.

A research video by MIT once showed a group of researchers having a meeting around the table, and the room was “listening in” on the conversation going on. Whenever a certain point was reached, such as deciding to add a new item to the agenda, or delegating a task to a member in the room, everybody had to shut up, and the moderator would speak the corresponding commands for the computer to keep up with what was going on. It was the worst group support interface imaginable.

Good group support software follows what’s going on in the room as good as it can, trying to detect from a variety of sensors, models, and other input what the current activity and actors are, and then takes initiative on a simple, reliable level to help the actors, without presuming to understand more than it can.

Computer scientists will argue that deriving this information from sensor values is not reliable, so the computer needs clear commands in order not to do something wrong. This is perfectly true in distributed settings with low bandwidth for human communication: If user A decides to pass control over the

shared mouse cursor to remote user B in a shared application, he usually has to click a button to do so.


In a collocated setting of an AE, an enormous advantage comes to the help of the system: social protocol. The people in the room can see and hear each other. If one person is controlling the mouse cursor using their laptop, and someone else wants to take over with their own laptop, they will just say so. The computer does not need to understand this verbal command, nor does he need to lock the cursor for everybody else but one user at a time: It can simply accept cursor movement from everybody in the room; if there's a conflict of concurrent access, the users will quickly and easily notice and resolve it among themselves. This approach, on the other hand, saves the users having to send explicit messages each time they wish to pass control of that cursor to someone else, making the interaction much more fluid.

Examples include the design of the interaction for the iRoom's remote cursor control that allows "mouse fights" to occur, simply always using the last coordinate received; or its iClipboard feature that lets people cut and paste in a single shared clipboard for the room.

Winograd et al., in their chapter elsewhere in this book, reflect on this concept by suggesting room infrastructure in which "...users and social conventions in an environment take responsibility for actions, and the system infrastructure is responsible for providing a fluid means to execute those actions."

Therefore:

Do not put unnecessary protocols into place that are aimed at avoiding overlapping access to technology, if that collision can be easily noticed and fixed by the users through social interaction. If a user issues a social protocol act, such as passing a wireless mouse to someone else, never require an additional repetitive step from the user to tell the room what he just did for everyone else to clearly see.



This is a basic pattern with no further references within this language.

17 EVENT NETWORKING **

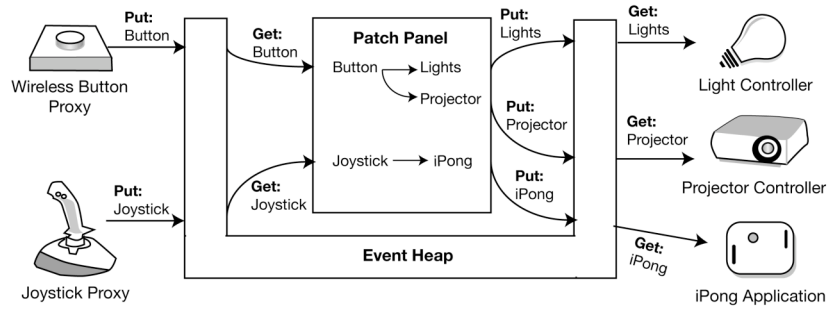


Fig. 18: The flow of events from input devices and applications, through a central networked event management (EventHeap) and mapping (PatchPanel) service, to the appropriate output applications and devices.

...you have installed your hardware and made OS choices. Now it is time to think about the middleware to let your users collaborate fluidly in the AE.



It is easy to link computers over a data network and send data using application-specific protocols. However, in an AE the machines, operating systems and applications must exchange information about events occurring on them in a more general and useful way.

One of the key principles in any operating system that supports interactive applications is *event processing*—button clicks on a mouse are passed via drivers into system-wide queues, where they are ordered, prioritized, and then distributed to their target application, typically depending on which application has the input focus at the moment. This mechanism has been highly optimized in modern operating systems, to ensure immediate feedback and continuous control for the user.


AEs require a similar mechanism, but now at the level of communication across multiple computers and operating systems. This means that everything needs to be internet-controllable, from devices such as projectors, speakers, and lamps, to applications such as presentation and brainstorming apps. For example, the Stanford iRoom and KTH iLounge were equipped with projectors that could be turned on and off via web-based interfaces. Scripting interfaces for applications such as PowerPoint were used to provide remote control of presentations.

The standard will most likely be a TCP/IP-based protocol, usually via a web interface. But it can also be another technology such as Bluetooth, or a mix of

these. The important thing is that all components in the room support being set up and controlled via events.

Therefore:

Provide each technology component in your AE with a control interface using an open networking standard.



Once devices and applications can be controlled via events, you can create services based on that infrastructure—COLLOCATED GROUP SERVICES (18) ...

18 COLLOCATED GROUP SERVICES **




Fig. 19: Users collaborating in a TeamSpot™. (Image copyright 2004–2008, TideBreak™, Inc. All rights reserved.)

...you have found a way to send information about events between your room components in a meaningful and practical way—EVENT NETWORKING (17), and prepared the room to share media between participants—INTERACTIVE WALL (8), AUDIO SPACE (14). Now you need to decide on what services to build on top of this infrastructure.



Today's operating systems and applications make working alone quite effective, but in an AE the group will want to work on information together, with frequent exchanges of digital information between the participants—and this is badly supported in today's standalone systems.

The Stanford Interactive Workspaces project tackled this problem by developing a robust, fault-tolerant middleware that built on the concept of the Event Heap, a tuplespace with publish-subscribe semantics (see [Borchers, Ringel, Tyler and Fox 2002] or Winograd's chapter for details). It allowed applications to send or subscribe to particular events, enabling applications, devices and services to talk to each other. Its loose coupling allowed individual components to disconnect without disrupting the remaining network.

On top of this, the software provided easy ways for room users to move their mouse cursor from their own laptop up to a shared screen in the room, with keyboard input following along. The package also enabled quick and easy sharing of URLs and arbitrary files or text and graphics clippings between people in the room that had the client software installed. An important principle of these tools promoted in part by the author was the use of core features such as copying and


pasting between computers without having to remember any new commands or keyboard shortcuts, instead simply using the existing commands and shortcuts for the expanded functionality.

The software was successful enough to be turned into a commercial product, *TeamSpot*, based on the same principles, that has since added screen sharing, full cross-platform support, and many additional features and improvements to the mix.

Since then, other software packages have developed similar ideas, though usually not as comprehensive. One example is DropCopy (<http://10base-t.com/software/macintosh/dropcopy.html>) for quick sharing of files between Macs and iPhones.

Therefore:

Provide seamless means for users of your AE to move files, text snippets, graphics and other media, URLs, and entire screen contents around the room to other personal laptops or shared displays.



This is a basic pattern with no further references within this language.

19 OVERFACE **



Fig. 20: In the i-LAND® environment, windows are crossing screens and machines, and users work with a seamless interface [Streitz, Tandler, Müller-Tomfelde, Konomi 2001].

...you have chosen what devices and services to make available in your AE, and are trying to come up with a good user interface representation to control everything—ROOM CONTROLLER (15).




People working in an AE tend to think of the environment and its computing services as a single, coherent entity. However, today's Personal Computers and their operating systems and applications are designed around a single-person interaction model.

The iRoom provided a web-based interface to control the entire room, and allowed for task-centered actions such as “make the room ready for a meeting” to be set up as commands, which would have turned on lights and projectors, launched certain applications, and brought up appropriate work documents and web pages.

At one point, we created a “Start the Room” button using our iStuff toolkit [Borchers, Ringel, Tyler and Fox 2002] that could simply be tacked to a wall near the entrance and would complete this setup procedure when pressed.

Therefore:

Give the user the illusion that he is interacting with a single, coherent user interface, not with an assortment of individual systems with different rules, metaphors and conceptual models.



This is a basic pattern with no further references within this language.

7 About the author

Jan Borchers is a professor of computer science and head of the Media Computing Group at RWTH Aachen University in Germany. With his group, he explores human-computer interaction beyond the desktop, from interactive workspaces and ubiquitous computing environments, to audio and video streams, to mobile and wearable devices. He has been involved in the research and design of interactive workspaces since 1995, including the Classroom of the Future project for the Ars Electronica Center in Linz, Austria [Mühlhäuser et.al. 1996], the Stanford Interactive Workspaces (iRoom) project [Borchers, Ringel, Tyler and Fox 2002] and its middleware software later commercialized as *TeamSpot* (see also the chapter by Johanson, Fox and Winograd), the interactive workspaces at Stanford's Wallenberg Hall and the iLounge at KTH (see also the chapter by Jansson), and lately the Aachen Media Space [Borchers 2006] at RWTH Aachen University, the author's present institution, as well as three similar spaces for *HumTec*, an interdisciplinary research center on humans and technology at the same university (<http://humtec.rwth-aachen.de>).

8 Acknowledgements

The author would like to thank the members of the RUFÄE team for the many shared projects, discussions and demonstrations that led to this pattern language, and Thorsten Karrer and the other members of the Media Computing Group for their support in collecting many of the photographs in this language.

This work was funded in part by the German B-IT Foundation, and by the German Government through its UMIC Excellence Cluster for Ultra-High Speed Mobile Information and Communication at RWTH Aachen University.

References

- Christopher Alexander, Sara Ishikawa, Murray Silverstein, Max Jacobson, Ingrid Fiksdahl-King, and Shlomo Angel. *A Pattern Language: Towns, Buildings, Construction*. Oxford University Press, 1977.
- Christopher Alexander. *The Timeless Way of Building*. Oxford University Press, 1979.
- Christopher Alexander. Keynote Speech, OOPSLA'96 11th Annual ACM Conference on Object-Oriented Programming Systems, Languages and Applications (October 6-10, 1996, San Jose, California), 1996. (Conference video).
- Apple Computer. *Macintosh Human Interface Guidelines*. Addison-Wesley, 1992.
- Christopher Alexander, Murray Silverstein, Shlomo Angel, Sara Ishikawa, and Denny Abrams. *The Oregon Experiment*. Oxford University Press, 1988.

- Elisabeth Bayle, Rachel Bellamy, George Casaday, Thomas Erickson, Sally Fincher, Beki Grinter, Ben Gross, Diane Lehder, Hans Marmolin, Brian Moore, Colin Potts, Grant Skousen, and John Thomas. Putting it all together: Towards a pattern language for interaction design. *SIGCHI Bulletin*, 30(1):17–23, January 1998.
- Kent Beck and Ward Cunningham. Using pattern languages for object-oriented programs. Technical Report CR-87-43, Tektronix, Inc., September 17, 1987. Presented at the OOPSLA'87 workshop on Specification and Design for Object-Oriented Programming.
- Jan O. Borchers, Richard N. Griffiths, Lyn Pemberton, and Adam Stork. Pattern languages for interaction design: Building momentum. Workshop at CHI 2000 (The Hague, Netherlands, April 2–3, 2000); report to be published, 2000.
- Jan Borchers, Eric Lee, Wolfgang Samminger, and Max Mühlhäuser. Personal Orchestra: A Real-Time Audio/Video System for Interactive Conducting. *ACM Multimedia Systems Journal Special Issue on Multimedia Software Engineering*, 9(5):458–465, March 2004.
- Jan Borchers. *A Pattern Approach to Interaction Design*. Wiley Series in Software Design Patterns. Wiley, 2000.
- Jan O. Borchers. CHI meets PLoP: An Interaction Patterns Workshop (at ChiliPLoP'99 Conf. on Pattern Languages for Programming, Wickenburg, AZ, March 16–19, 1999). *SIGCHI Bulletin*, 32(1): 9–12, January 2000.
- Jan O. Borchers. Interaction design patterns: Twelve theses, 2000. Position paper for the workshop *Pattern Languages for Interaction Design: Building Momentum*, CHI 2000 (The Hague, Netherlands, April 2–3, 2000).
- Jan Borchers. Teaching HCI patterns: Experience from two university courses. In Position Paper for the HCI Patterns Workshop at the CHI 2002 International Conference on Human Factors and Computing Systems. ACM, April 2002.
- Jan Borchers. The Aachen Media Space: Multiple displays in collaborative interactive environments. In *CHI 2006: Workshop "Information Visualization and Interaction Techniques for Collaboration across Multiple Displays"*, Montreal, Canada, April 2006.
- Jan Borchers, Meredith Ringel, Joshua Tyler, and Armando Fox. Stanford interactive workspaces: A framework for physical and graphical user interface prototyping. *IEEE Wireless Communications*, 9(6):64–69, December 2002.
- Lon Barfield, Willie van Burgsteden, Ruud Lanfermeijer, Bert Mulder, Jurriënne Ossewold, Dick Rijken, and Philippe Wegner. Interaction design at the Utrecht School of the Arts. *SIGCHI Bulletin*, 26(3):49–79, 1994.
- Andrew Dearden and Janet E. Finlay. Patterns in HCI: A critical review. *Human Computer Interaction Journal*, 21(1):49–102, 2006.
- Satinder P. Gill and Jan Borchers. Knowledge in co-action: social intelligence in collaborative design activity. *AI & Society Journal of Human-Centred Systems and Machine Intelligence*, 17:322–339, October 2003.
- Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley, 1995.
- Åsa Granlund and Daniel Lafrenière. A pattern-supported approach to the user interface design process. Workshop report, UPA'99 Usability Professionals' Association Conf. (Scottsdale, AZ, June 29–July 2, 1999), http://www.gespro.com/lafrenid/Workshop_Report.pdf, 1999.
- François Guimbretière, Maureen Stone, and Terry Winograd. Fluid interaction with high-resolution wall-size displays. In *UIST '01: Proceedings of the 14th Annual ACM Symposium on User Interface Software and Technology*, pages 21–30, New York, NY, USA, 2001. ACM.

- Jefferson Y. Han. Low-cost multi-touch sensing through frustrated total internal reflection. In *UIST '05: Proceedings of the 18th annual ACM symposium on User interface software and technology*, pages 115–118, New York, NY, USA, 2005. ACM.
- Steve Hodges, Shahram Izadi, Alex Butler, Alban Rrustemi, and Bill Buxton. Thinsight: versatile multi-touch sensing for thin form-factor displays. In *UIST '07: Proceedings of the 20th annual ACM symposium on User interface software and technology*, pages 259–268, New York, NY, USA, 2007. ACM.
- Max Mühlhäuser, Jan Borchers, Chris Falkowski, and Knut Manske. The conference/classroom of the future: An interdisciplinary approach. In *Proceedings of the IFIP Conference on the International Office of the Future: Design Options and Solution Strategies*, pages 219–235, Tucson, Arizona, April 1996. Chapman & Hall, London/New York.
- George A. Miller. The magical number Seven, plus or minus two: Some limits on our capacity for processing information. *The Psychological Review*, 63:81–97, 1956.
<http://www.well.com/user/smalin/miller.html>.
- Donald A. Norman and Stephen W. Draper. *User-Centered System Design: New Perspectives on Human-Computer Interaction*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1986.
- Donald A. Norman. *The Psychology of Everyday Things*. Basic Books, New York, 1988.
- Norbert A. Streitz, Jörg Geißler, Torsten Holmer, Shin'ichi Konomi, Christian Müller-Tomfelde, Wolfgang Reischl, Petra Rexroth, Peter Seitz, and Ralf Steinmetz. i-land: an interactive landscape for creativity and innovation. In *CHI '99: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 120–127, New York, NY, USA, 1999. ACM.
- Norbert Streitz, Carsten Magerkurth, Thorsten Prante, and Carsten Röcker. From information design to experience design: smart artefacts and the disappearing computer. *interactions*, 12(4):21–25, 2005.
- Norbert Streitz, Peter Tandler, Christian Müller-Tomfelde, Shin'ichi Konomi. Roomware: Towards the Next Generation of Human-Computer Interaction based on an Integrated Design of Real and Virtual Worlds. In: J. Carroll (Ed.): *Human-Computer Interaction in the New Millenium*, Addison-Wesley, 2001. pp. 553-578.
- Chia Shen, Frédéric D. Vernier, Clifton Forlines, and Meredith Ringel. Diamondspin: an extensible toolkit for around-the-table interaction. In *CHI '04: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 167–174, New York, NY, USA, 2004. ACM.
- Jenifer Tidwell. *Designing Interfaces: Patterns for Effective Interaction Design*. O'Reilly, 2005.
- K. van Duyne, James A. Landay, and Jason I. Hong. *The Design of Sites: Patterns for Creating Winning Web Sites*. Prentice-Hall, 2 edition, 2006.
- Stefan Werner. An algorithm for audio skew compensation in low latency environments. In *Proceedings of the ICMC 2005 International Conference on Computer Music*, San Francisco, CA, 2005. International Computer Music Association (ICMA).