

Breaking the Screen Barrier

Lars Erik Holmquist

1 Introduction

In a video presentation devised in 1981, Robert Spence and Mark Apperley presented a vision of the future office environment [38]. They saw the office of the future as a place for rich, multi-modal interaction with digital information, where users had access to a variety of input and output methods – displays of a variety of sizes, from desktop to wall-sized; gestural interaction; voice input; handwriting recognition; and so on. This was a compelling view, where users would take advantage of their whole environment to access and manage digital information, rather than being limited to a small single screen and the “point-and-click” interaction we are used to today. Although several other researchers were exploring similar avenues (e.g. the “Put-that-there” system at MIT’s Architectural Machine group [7]), it is safe to say that Spence and Apperley’s vision was presented well before its time. Even today, this kind of rich mix of digital and physical space is far away from being widely implemented, and is still very much at the stage of research prototypes (see e.g. [29]).

However, the reasons for why our current interaction with digital information is not as rich as that envisioned by Spence and Apperley is worth some thought. Certainly, the technology needed is quite complex – but on the other hand the authors asserted that most of their vision could be implemented with then-current technology. A more interesting “obstacle” can instead be found in what is arguably the most successful innovation in the human-computer interaction (HCI) field so far: the *Windows-Icons-Menus-Pointers* (WIMP) interaction paradigm, and the related notion of a *Graphical User Interface* (GUI), where the digital information is presented according to a *desktop metaphor*. Developed at Xerox PARC in the 1970’s, and first introduced commercially with the Xerox STAR computer in 1980 [20], WIMP and GUI did not become a commercial success until the Macintosh computer was introduced by

Apple in 1984 [1], with Microsoft Windows becoming the dominant GUI several years later [12].

WIMP, GUI and the desktop metaphor has now become the totally dominant mode for interacting with computers. Most observers would agree that this is because these were very brilliant ideas. However, a problem with ideas that are as brilliant as these is that they can become very hard to look beyond, and there is a growing realization of this fact within the human-computer interaction (HCI) research community [9, 18]. In recent years we have seen an increasingly intensive search for new alternatives. There have of course been a large number of attempts to find ways to enhance user interaction while staying within the limits of WIMP and GUI. Several alternatives to the desktop metaphor have been introduced, for instance by using time rather than space as an organizing mechanism [16, 31], or by the introduction of a 3D-graphics element to take better advantage of the user's spatial perception [33]. The limited amount of space available on a desktop screen has been addressed in a number of ways, for instance by the introduction of several separate workspaces [19] or by a variety of information visualization techniques that present data more effectively [10]. Of particular interest for this thesis are those information visualization methods which attempt to "expand" the area available for showing information, through the introduction of visual distortion; notable examples include the Bi-Focal display (which was part of Spence and Apperley's vision of the future office) [38, 39] and the Perspective Wall [26].

In the years in which WIMP and GUI came to dominate the commercial side of human-computer interaction, researchers also introduced an ever-growing set of divergent approaches. *Virtual reality* promised to place users inside in a virtual world of information, where interaction would be as rich as or richer than in the "real" world [32]. *Ubiquitous computing* was based on the notion that computers would become available everywhere. It aimed to move computer usage away from the desktop-centric WIMP approach and into the user's physical environment, through the introduction of computers in various size and shapes that were specialized for different usage situations [42]. *Augmented reality* proposed to mix computational properties with real-world objects, either through a projected graphical overlay, wearable see-through computer displays, or some other computational augmentation of real-world objects [43]. *Intelligent environments* were proposed as physical environments where a variety of sensors, cameras, etc. would watch the user

and where the technology embedded in the surroundings would respond according to the user's explicit or implicit wishes [11]. *Wearable computers* aimed to take the computer away from the environment and instead let it become like a piece of electronic clothing that would always be "on" and support the user at all times [27]. *Graspable interfaces* proposed providing certain functions of the GUI as physical instantiations for a more direct physical interaction [17], and in an extension of this, *tangible media* aimed to remove the border between the physical and digital world altogether, in order to "change the world itself into an interface" [21].

But however compelling these alternatives may be, there is little doubt that for many tasks, WIMP and GUI are the best modes of interaction currently available. Since they were developed specifically with the office worker in mind, typical office applications such as word processors and spreadsheets are hard to imagine functioning outside the confines of the desktop computer. Mobile technology may offer us the promise of accessing our digital information any time, anywhere, but it will probably be long before most of us can conveniently write longer texts or manage complex calculations on mobile phones or PDAs. That said, many of the uses for computers are now completely outside the office domain, and with computer technology becoming ever smaller and easier to integrate in our daily life, we will see computer technology appear in many situations we have not even dreamed of. Thus, WIMP and GUI will probably stay as the dominant interaction paradigm for most of the tasks we have traditionally associated with computer use, but there will be a myriad of complementary ways in which computers appear in our lives, and these will require radically different approaches to interaction.

This thesis spans a large part of the spectrum of interface approaches outlined above, from interaction strictly within the WIMP and GUI domain, over hand-held GUI computers, all the way to wearable computers and tangible media. The common thread is an effort to free the user from the inherent limitations of the computer screen – to "break the screen barrier", figuratively or literally. Despite the progress in recent years in both size and resolution, the typical user is still limited to a screen 17 to 19 inches in size, and with a resolution of no more than approximately 1000 by 1000 pixels. Furthermore, a computer screen is usually fixed to the same location, its mode of presentation is inherently 2-dimensional ("flat"), and it offers no provision for tactile feedback or

the use of any other sensory modalities than sight. Although the computer screen is an incredibly flexible and powerful canvas for interaction design, it is clear that it is also limited in many ways. This thesis presents a few possible strategies for overcoming some of these limitations.

2 The thesis: Breaking the screen barrier

The thesis consists of five papers, four of which were published in 1999 and one which has not yet been published, but which is based in part on publications from 1997 and 1998. Apart from the required reformatting to fit the format of the thesis, the published papers are presented in unaltered form. The papers are as follows (when Holmquist is not listed as first author, authors have been listed alphabetically and/or according to affiliation):

1. Holmquist, L.E.: *Flip Zooming: Focus+Context Visualization of Linearly Ordered Discrete Visual Structures*.

Submitted for publication.

Based in part on the following short papers:

Holmquist, L.E.: Focus+Context Visualization with Flip Zooming and the Zoom Browser. In *Extended Abstracts of ACM SIGCHI Conference on Human Factors in Computing Systems (CHI '97)*, pp. 263-264, ACM Press, 1997.

Holmquist, L.E. and Ahlberg, C.: Flip Zooming: A Practical Focus+Context Approach to Visualizing Large Information Sets. In Smith, M.J., Salvendy, G. and Koubek, R.J. (eds.), *Design of Computing Systems: Social and Ergonomics Considerations (HCII '97)*, pp. 763-766, Elsevier Science B.V., 1997.

Holmquist, L.E. and Björk, S.: A Hierarchical Focus+Context Method for Image Browsing. In *Computer Graphics Annual Conference Series Abstracts and Applications (SIGGRAPH '98)*, p. 282, ACM Press, 1998.

Björk, S. and Holmquist, L.E.: Formative Evaluation of a Focus+Context Visualization Technique. Poster presented at *Annual Conference of the British HCI Group (HCI '98)*, Sheffield, UK, 1998.

2. Björk, S. Holmquist, L.E. and Redström, J.: *A Framework for Focus+Context Visualization*.

Abridged version in *Proceedings of IEEE Symposium on Information Visualization (InfoVis '99)*, pp. 53-56, IEEE Press, 1999. Full version in *CD-ROM Proceedings of IEEE Visualization 1999*, IEEE Press, 1999.

3. Björk, S., Holmquist, L.E., Redström, J., Bretan, I., Danielsson, R., Karlgren, J. and Franzén, K.: *WEST: A Web Browser for Small Terminals*.

In *CHI Letters Vol 1 Issue1, Proceedings of ACM CHI Conference on User Interface Software and Technology (UIST '99)*, pp. 187-196, ACM Press, 1999.

4. Holmquist, L.E., Falk, J. and Wigström, J.: *Supporting Group Collaboration with Inter-Personal Awareness Devices*.

Journal of Personal Technologies, 3(1-2), pp. 13-21, Springer Verlag, 1999.

5. Holmquist, L.E., Redström, J. and Ljungstrand, P.: *Token-Based Access to Digital Information*.

In *Proceedings of First International Symposium on Handheld and Ubiquitous Computing (HUC '99)*, pp. 234-245, Springer Verlag, 1999.

The theme of the thesis is “breaking the screen barrier”, i.e. to somehow overcome the inherent limitations of the computer screen. The work started out within the WIMP / GUI domain with the intention of expanding the workspace available to a computer user, through the introduction of a means to show more information on a desktop computer screen. This resulted in the so-called *flip zooming* visualization technique presented in the first paper of the thesis, *Flip Zooming: Focus+Context Visualization of Linearly Ordered Discrete Visual Structures*. The technique could be used to visualize documents, images, and other data sets. Flip zooming was designed with a traditional desktop computer in mind, but represents a quite different approach compared to windowing systems and the desktop metaphor. By virtually expanding the available workspace to become much larger than the physical screen, it was a first step towards “breaking the screen barrier”. Subsequent work with flip zooming resulted in the more general framework for focus+context visualization presented in the second paper, *A Framework for Focus+Context Visualization*, where a formal method for describing and constructing focus+context visualizations was introduced. This work,

though based in the GUI domain, is possible to generalize to other interaction approaches, e.g. 3D-environments.

Furthermore, we soon realized that the flip zooming technique could be applied to other devices apart from WIMP computers, such as mobile computers with wireless internet connections. This resulted in *WEST: A Web Browser for Small Terminals*, which is also the title of the third paper. Although GUI based, the input and output capabilities as well as the usage conditions of a mobile terminal are drastically different from that of a desktop computer. This proved to require some major changes in our approach to interface design, which are partly documented in the paper and which we have continued to explore in later works (e.g. [3]). Mobile computing, whether GUI-based or built on other approaches, is likely to be a major area for future HCI research.

In parallel with the work of flip zooming, we also explored other, quite different, approaches to human-computer interaction. *The Hummingbird* was a specialized wearable computer, which did not in any way resemble a desktop computer, and it had no GUI whatsoever. The Hummingbird and some preliminary evaluations are discussed in the fourth paper, *Supporting Group Collaboration with Inter-Personal Awareness Devices*. The purpose of the Hummingbird was quite similar to many desktop-based awareness systems, but by being completely mobile and independent of any infrastructure, it acknowledged the fact that users spend much of their time away from the desktop. Thus Hummingbirds literally “broke the screen barrier” by moving the interaction away from the desktop computer completely. This work has continued with the development of a second generation of Hummingbird prototypes and a variety of evaluations (e.g. [41]).

Finally, our work with the *WebStickers* system [24] led to the fifth and final paper. WebStickers let users couple everyday physical objects with digital information, in a fashion similar to some systems for augmented reality. In the paper, *Token-Based Access to Digital Information*, we used the lessons learned from the WebStickers system to draw some general conclusions about proposed interaction paradigms such as tangible media and graspable interfaces. This paper generalizes the interaction with computers to such an extent that screen output is just one of a variety of possible interaction methods, so that the user’s whole physical space here becomes an arena for accessing and ordering digital information – the “screen barrier” has truly been broken.

Taken together, the papers also represent a growing realization that the notion of “context” as it has been used in research on graphical user interfaces is very limited, since it only acknowledges interaction taking place on a computer screen. Instead, “context” in human-computer interaction should in reality also include the whole physical world within which the user interacts. This is consistent with the many proposed alternatives to WIMP and GUI that were outlined in the introduction. At the same time, the thesis also reflects the fact that WIMP and GUI will be with us for yet some time and that it is still very worthwhile to explore enhancements and augmentations within that domain, a notion that is sometimes easy to forget in the rush towards the various tangible, mobile, ubiquitous and wearable alternatives. By playing off different interaction paradigms ranging from WIMP and GUI to augmented reality, tangible interfaces and wearable computing, the papers in the thesis thus explore how human-computer interaction can “break the screen barrier” both in a figurative and a literal sense.

3 Research method

This thesis is a work in the Swedish scientific discipline called *informatics*, or more specifically the *new informatics* as proposed by Bo Dahlbom in 1996 [13]. The Swedish discipline of informatics could be described as “applied computer science”, and has its roots in information systems research. The subject matter of the new informatics is *information technology (IT) use*. However, the new informatics is not focused only on studying the use of IT – it is also very much interested in changing and improving the use of information technology. In other words, it is a *design oriented discipline* [13, p. 29]. According to Dahlbom, “whatever we do with our discipline (...) we should protect our design interest” [13, p. 30], stressing the importance of the discipline’s active involvement in and contribution to the development of information technology and its use. This sentiment is in line with the work in this thesis, which has been very much concerned with the development of novel IT artifacts and the exploration of their use.

The new informatics’ strong interest in design follows from the notion that it is an *artificial science*, i.e. one that is concerned with objects created by man, as opposed to the *natural sciences*, which study

things in nature. The idea of an artificial science was presented by Herbert Simon in a series of lectures in 1968, which were published in book form in 1969 [34], and revised several times, most recently in 1996 [35]. Simon argued that there is a need to acknowledge the importance of design in the engineering disciplines, rather than to have them fall into the trap of mimicking the methodology of the natural sciences in an effort to gain scientific credibility. Influenced by then-current research in artificial intelligence, Simon also proposed a theory of design, where the design activity takes the form of a search for the best among a certain set of available alternatives. However, a major problem seems to be that Simon's design theory is unable to capture the creative, intuitive and accidental elements in the design process, which often can mean the difference between a brilliant design and one that is merely satisfactory.

Dahlbom has addressed part of this problem in a work that is a direct response to Simon [14]. Although an enthusiastic supporter of the general idea of artificial sciences, Dahlbom presented several objections to Simon, having mainly to do with Simon's definition and theory of design and its role in the artificial sciences [14, pp. 5-6]. Dahlbom then proceeds to propose a set of foundations for an artificial science, four of which are directly inspired by Simon's work, and four which also are inspired by Simon but are more specifically based on the foundations of the natural sciences. I will summarize the eight points below; for more details see [14, pp. 6-9]. (Dahlbom's text is in italics; the summaries are mine.)

1. *Artifacts are designed rather than described*, i.e. artificial science studies possibilities rather than the already realized; also, it insists on concrete realization as a way to make sure something is really possible.
2. *Technology in use rather than design practice*, i.e. the research will be design oriented from a user perspective, considering improvements in use quality rather than product quality per se.
3. *Artifacts have quality rather than functionality*, i.e. we will have to go beyond thinking of the relations between people and technology just in terms of "use", and introduce dimensions such as aesthetics, symbolism, ethics and politics.
4. *Artificial science is normative rather than objective*, i.e. since artificial science is concerned with more than an artifact's functionality, it

involves the idea of the “good life”, and thus it goes beyond purely objective concerns.

5. *Artifacts are accidental rather than essential*, i.e. rather than being nicely broken down into simple principles, artificial laws are local design solutions rather than general principles, and the artificial world is haphazard and provisional.
6. *Artifacts are constructed rather than documented*, i.e. results are only judged on the basis of successful construction – it is the quality of the technology that matters, not the documentation.
7. *Artificial science has heuristics rather than methods*, i.e. whereas the emphasis on methods turn natural science into a bureaucratic administration of ideas, creativity takes the center stage in artificial science, requiring a reliance on heuristic rules of thumb, intuition and tacit knowledge, experience and tinkering.
8. *Artificial science is engaged rather than disinterested*, i.e. the artificial sciences are not interested in objective detachment but in interacting with artifacts, and values play an important role in how one chooses what to work with, making artificial sciences more openly politicized.

The details of these points should of course be open for debate and discussion, but they are all quite easy to accept for this author, and as we shall see we will have no problem placing the work presented in this thesis within this framework. But what about the *process* of designing new artifacts? If we want to produce innovations in the field of information technology and IT use, what is the process we should follow? Dahlbom says that “artificial science is not a theoretical study of the design of concrete artifacts, but a systematic, institutionalized form of such design activity with the ambition to improve the world of artifacts” [14, p. 5]. Several of Dahlbom’s principles stress the importance of practical and use-oriented design, but fail to give readers much help in how this design actually happens.

Going back to Simon, design is “... concerned with how things ought to be, with devising artifacts to attain goals.” ([35, p. 114]). This sounds very simple: we take a look at the world we live in, we figure out how it might become better, and we design an artifact which takes us to this new, better world. There are indeed good methods to help us understand the world we live in; ethnographical studies, economical models, etc. There is also an abundance of information technology that could help us

in attaining our goals – not just desktop computers, but also mobile devices, embedded chips, smart materials, sensors and actuators, and so on. Thus, the stage seems well set for the new informatics researcher: armed with a well-founded knowledge of the world and a set of novel information technology, he or she can easily proceed to change the world into a better place!

But as the reader probably already has noticed, there is one thing missing: the goal, that better place the world should be. How we find out what this goal is seems to be the one crucial thing missing from Simon's definition of design. Dahlbom addresses that this is a problem, in particular by bringing in other dimensions such as aesthetics and politics, thus implying that we should strive for goals that increases the "happiness" of the user, but of course in the same breath acknowledging that happiness is a very much a "rubber concept" [14, p. 7]. This does not really take us very much further: yes, we should strive to find "good" goals, but what are they and how do we find them?

One approach can be found within the first strand of the new informatics to define its own research agenda: *mobile informatics* [23]. Mobile informatics is aimed at inventing new IT use in mobile settings through interdisciplinary collaboration, and identifies two methods for idea generation (i.e. goal definition): idea generation informed by studies of current practice; and technologically informed idea generation. But again, this does not help us in understanding the design process – that we know what the world is like, and what technology can currently do, does not guarantee that we will come up with innovative (let alone "good") uses for technology. In fact, this is still very close to Simon's definition of design (above). However, a hint of what the process might entail is found in the emphasis on practical implementation of artifacts, which is identified as being necessary for finding limitations and possibilities: "The very construction of the IT artifacts will (...) give rise to new insights (and) form the IT artifact being implemented" [23, p. 207].

An approach which has been described in detail within the framework of mobile informatics, is that of so-called *scalability through cultivation* [2]. Here, the goal is not to radically change a work situation through the introduction of novel IT use, but to support the current work practice through "guided evolution", addressing only the badly functioning parts. In practice, this is carried out by first doing a study of a workplace (inspired by ethnographical methods), and then devising design proposals based on the results of that study. The design proposals

are then presented to the workers for comments. This is reminiscent to the method known as *participatory design*, where the prospective users are directly involved in the construction of computer systems [28].

The case at hand concerned an order packaging department, but the methodology should be possible to generalize to many other situations. The most obvious strength in this approach is that since the resulting design proposal is based on a study of the workplace, rather than being a “flash of inspiration” thought up in isolation from the actual work practice, it should have a greater chance of addressing the important problems. Also, since the approach stresses the importance of “cultivation” rather than radical change, proposed designs should have a greater chance of being implemented without disrupting current work practice.

However, the fact that a proposal is based on a study does not guarantee that it is the best possible proposal, not even that it is a good one! The study and design proposals given in the paper can be taken as an example. For instance, the study shows that physical work orders (pieces of paper) that are taken from a communal notice board are used as a coordination mechanism. The proposed design solution *removes* the physical work orders and replaces them with a system consisting of a large computer screen (replacing the notice board) and a set of networked hand-held computers (containing the work orders), one for each worker. The authors claim that using work orders on paper is less efficient than using computers, but in fact there are several studies that show how physical artifacts such as paper are an important support in many work processes, even those which rely heavily on computer support (e.g. [25]). Although removing such a mechanism and replacing it with computer technology might in some cases make work more effective, it might just as well spell disaster in the current work practice! Until the proposed design has been implemented it is impossible to know what the effect might be. In any case, an alternative design proposal which integrated the existing coordination mechanism of the physical work order with some kind of computer support might have presented a better solution to many of the problems identified in the study.

The reason for bringing up this example is not to criticize it in depth, but to point out that studies, no matter how well done, are always very much open to interpretation. Although the facts may be correct, this does not guarantee that the conclusions are. Even more importantly, the quality of design proposals that are based on such a study is not in any

way certain; one person might have a brilliant idea based on a study, whereas another might come up with something which does not improve the situation in any way, and perhaps even makes it worse. The only way to know is to carry through with the “cultivation” – to really implement the proposal and see what happens, and then make further adjustments based on that. Thus, until a proposal resulting from an observation has been taken back to the workplace and implemented, there is really no way of judging its quality; just asking the workers if they like the proposal will certainly not be enough.

Another objection to the cultivation approach is that it only promotes incremental change, not radical innovation. This is probably as it should be; if we want to cultivate a current work practice, and improve it without destroying it, cultivation is a reasonable approach. However, if we want to introduce a new work practice, or if we are interested in true innovation, cultivation might not work very well. Any radical departure from the current work practice will then seem as a threat and quite probably be met with opposition. Furthermore, many truly ground-breaking innovations – the telephone, the car, the internet, and so on – are not specific but general in their use and effects, and thus not rooted in any easily studied practice. In fact, it is quite hard to see how such innovations could arise from any kind of workplace study.

Returning to the work in this thesis, rather than starting with studies of work practice, it has been based on a practical approach to design, fitting well with Dahlbom’s foundations for an artificial science. We have focused on producing *innovations* rather than incremental improvements. But innovations have no value in themselves – they need to be proven useful; to be doable, to actually have a place in the real world. The only way of doing this is by practical implementation. By not only dreaming about innovative artifacts but actually *creating* them, we will find out much more about them than is possible by staying on a purely conceptual level. When an artifact is actually built and tried in the real world, rather than just presented as a design suggestion, one will often find it to be a very different beast than one thought. Not only does the physical implementation of artifacts tend to turn up many unexpected problems; more interesting for the new informatics researcher is that the use of a novel artifact is never uncomplicated. Unexpected things happen when people get their hands on artifacts – things the developer could never have foreseen. This element of uncertainty, the fact that when real people are allowed to use an artifact they will offer both criti-

cisms and suggestions, should be at the center of new informatics research. This response in turn should be channeled back into the design of the artifact, so that it may change accordingly.

Thus, the process of doing this research can not be broken down cleanly in “before-after” situations, or in “states” and “goals”, because the goals are not known – or at least not very well specified – when the work begins. When approached from such a practical perspective, the reason for the lack of a clear methodology to find design goals is probably this: the goal in any design process is very much a moving target, and design in itself is a crucial factor in making it move. Rather than being something which is defined at the start of a project and reached at the end, the goal of most design processes changes continually as the design evolves. This is especially true when design takes place within a research context, where the interest often lies as much or more in finding questions as in answering them. Thus, there is a need to complement the new informatics with a more detailed theory of the design process, which manages to take into account Dahlbom’s foundations of an artificial science.

The process of designing computer software might offer a useful parallel.¹ While programming originally was an informal practice undertaken by a single person or very small teams, it has now reached the stage where it is often a time-critical effort involving hundreds or even thousands of people. Several models have been proposed and used to steer the software development process. An early model was the *waterfall model* [5], where development is performed in a series of well-defined steps, where the software is “frozen” at each stage. This is very similar to well-functioning methods for developing complex hardware systems, where each component is specified and produced according to certain criteria; however, hardware is very different from software. In software, “once the drawings and models (programs) are complete, the final product exists” [4, p. 30]. An analogy in [4] to house building versus sculpting is enlightening: whereas a house is built with a good understanding of the requirements, and modifications are restricted to cosmetic and minor items, a sculpture is much less rigid, since clay can be added and subtracted during the whole process. In sculpting, the process of making the sculpture is part of finding out what the sculpture

1. The following two paragraphs on software development is based on [36]; quotations are not first-hand, but taken from that text.

will look like, which fits well with our notion of the design process in new informatics research.

Several models that try to capture this iterative nature of software design have been proposed, with Boehm's *spiral model* perhaps the most well known [6]. In this model, a software project will start with a set of requirements, but these are not cast in stone; instead, they will change and develop as the development continues. Through a series of risk analysis, prototyping, and verification cycles, a piece of working software is constructed. The difference between this and the waterfall model is that software is never "frozen"; instead, it evolves, with changes and amendments happening over time.

Another spiral model, which fits very well with the work in this thesis, is what we will term *Verplank's Spiral*, devised by William Verplank when working at Interval Research.¹ This model is not limited to software development but was developed to describe the general process of developing marketable products. The model stresses the flexible nature of design even more than Boehm's Spiral, since rather than starting out with a set of requirements, it starts with something as unspecific as a "hunch"! Also, whereas the models for software development are meant to steer or guide the work, Verplank's Spiral should be considered as more of a descriptive tool.

Verplank's model (**Figure 1**) is placed in a continuum where the vertical axis denotes the dimension of *paradigms* versus *industries*. In this model, a project typically starts out with a *hunch* – a vague notion of what to do. This leads to a *hack* – a first, primitive technical demonstrator of some kind. The hack makes it possible to *try* if the hunch is valid, and this in turn leads to an *idea*. The idea leads to one or more *designs*; this is the place in the spiral where several alternative avenues present themselves, since it might be possible to devise more than one design from the same idea. The designs are then fashioned into *prototypes* – working instantiations of the design. The prototypes can then be *tested*, and this in turn leads to a set of *principles* arising from the tests. These principles can then be fashioned into *plans*, which are specific enough to be used for *production* of actual products which reach the *market*.

1. The material on Verplank's Spiral is based on personal communication, William Verplank, Elsinore, Denmark, April 14, 2000. See also [15].

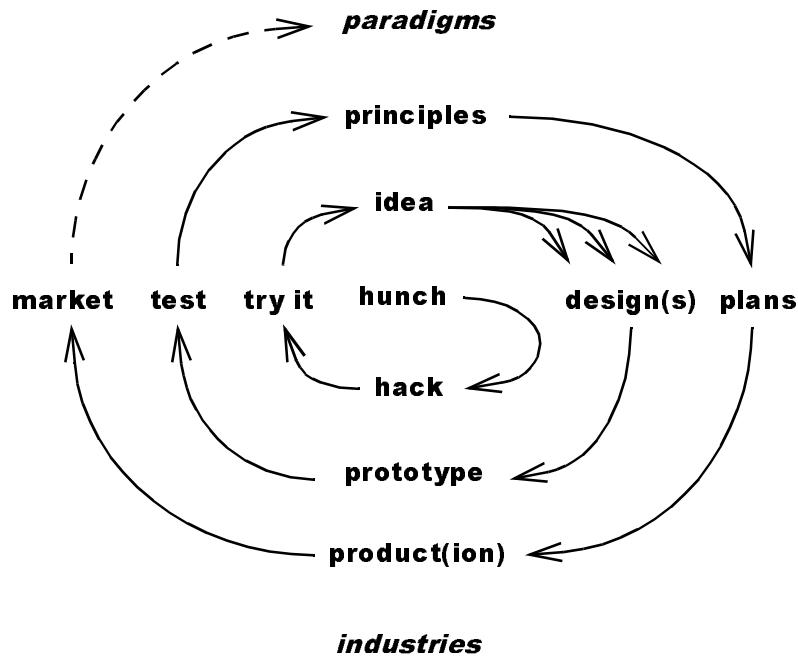


Fig. 1. Verplank's Spiral.

Finally, if the product is successful it might give rise to a new *paradigm*, that might even become an integral part of our lives (this would presumably happen quite rarely, and I have chosen a dashed line to indicate this, and also in acknowledgement that it is not absolutely clear to me whether Verplank actually meant paradigms to be an actual design goal as well as an underlying dimension). It is easy to see how all major technological inventions, from the printing press to aeroplanes, from television sets to the desktop computer, can be made to fit into this cycle.

The spiral is of most interest here because it allows us to find the place of various activities in research and development and to relate those to each other within the spiral. For instance, the so-called “demo-or-die” approach (most famously embodied by the MIT Media Lab [8]) is an example of an activity which stays very much within the “hunch” and “hack” domain. By creating an environment where hunches are encouraged and can be nurtured into hacks, so-called “demos”, such an institution can produce a wide variety of exciting new ideas which chal-

lunge the status quo. On the other hand, one might question the value of some of those ideas: if they do not reach further out into the spiral, it is hard to see how they might have a real impact on everyday life, especially if there are no resulting general principles to take away from them. But as long as the main purpose is to produce exciting innovations, having hunches and doing hacks is quite sufficient.

Similarly, an institution that concentrates on testing, such as for instance a usability lab, would be working very much in the domain of tests and coming up with general principles as a result. In this part of the spiral, the important thing is to examine existing prototypes and subject them to tests which are rigorous and wide-ranging enough to come up with general rules as to what works and what does not. This activity (as the spiral nicely illustrates) is much closer to developing plans that in turn can become actual products for the market, and might thus seem to be more relevant from the perspective of a commercial enterprise. On the other hand, by concentrating on evaluating and refining existing prototypes, the likelihood of producing ground-breaking innovations is probably much smaller than when working closer to the centre of the spiral.

Any truly successful design would thus start at the center of the spiral and move all the way out. But this is not the same as saying that the same person or team has to do everything! Although the occasional lone inventor might have been able to take a hunch all the way to a marketable product, any serious enterprise would certainly require several different specialists in different areas to go through with the whole process. This would seem to mean that there is nothing wrong with one group producing hacks, another doing designs and prototypes based on the resulting ideas, a third testing them and coming up with principles, etc. But having different parts of the design process separated also leads to problems, and decisions made early in the process might be hard to affect in the later stages, sometimes leading to products which are not as good as they should be. Donald Norman suggested *human-centred development* as a solution: by letting a user-centred perspective be a part of the process from the very start, he argues that better, more usable and less complex products will appear as a result [30]. This would presumably mean that the notion of a user would be present all the way from the first hunch, rather than being something which is added at the last minute when a product turns out to be too complex or user-unfriendly.

But it is also worth noting that Verplank's Spiral aims to describe the process of producing *products*; it is not certain that it is directly transferable to scientific research. Whereas a commercial company is mainly focused on producing artifacts that can be sold with a profit, a scientific institution is usually more interested in coming up with principles and paradigms. In fact, scientists might want to skip the product and market phases altogether and generate principles that are powerful enough to become paradigms in their own right. We can find many examples of this in the natural sciences. For instance, the "Big Bang" theory of the origin of the universe was certainly not developed with the intention of turning it into a successful product, but it has still entered into our consciousness and fundamentally altered the way we think about the world.

Can the same be said for the science of the artificial – can we cut products and markets out of the loop and still produce good scientific works? Probably not. A science of the artificial is concerned with the design of artifacts, and these artifacts are meaningless unless they enter into a relationship with a human being. If we define a "product" as an artifact which is produced with the intention that someone will want to use it, then artificial science is definitely concerned with products, whether we like it or not. The artifacts we produce may not be mass-market items; indeed, they might not satisfy a single person, not even the designer herself. But the perspective that someone, somewhere must interact with them seems to be fundamental to producing successful work in an artificial science, and thus also in the new informatics. This also seems to agree very well with Dahlbom's eight foundations of an artificial science. In the following we have therefore chosen to adopt Dahlbom's eight points as a general framework for the work in this thesis, and Verplank's Spiral as a tool to describe the type of work performed in the individual papers.

4 Relating the thesis to the research method

If we now return to Dahlbom's eight foundations, we can examine how well the thesis material fits within his definition of an artificial science. This work has relied on a practical approach to information technology, so that by designing and trying out novel designs, we have gained a better understanding of what is possible (point 1). It has been permeated by

a user-centred perspective, so that each artifact has been constructed with an idea of use and users rather than being examples of exciting and advanced technology (point 2). The resulting artifacts have been evaluated not according to how efficiently they solve a particular problem, but how the user experiences them in practice (point 3). Because of this, there has been no objective way of telling whether a design has been successful; rather, this has been done on the grounds of experience and user reactions (point 4). Although some of the artifacts were the results of well-thought-out research agendas, whereas others started as pure hunches, the process in arriving at the final results has in all cases been one of experimentation, staying with some solutions and discarding others (point 5). In the empirical work (papers one, three and four), every one of the artifacts has been a working instantiation of an idea; in the theoretical work (papers two and five) we have aimed to construct useful and productive frameworks rather than theory for its own sake (point 6). Creativity and accidents has been key ingredients in arriving at many of the final results, just like it is in all design work (point 7). And finally, the work has addressed key issues such as privacy, information overload, and how we can make information technology easier to use (point 8).

The general framework of an artificial science thus fits well with the work. What about the design process and Verplank's Spiral? Here, we must look at each paper individually. The first paper, *Flip Zooming: Focus+Context Visualization of Linearly Ordered Discrete Visual Structures*, started with the hunch that something could be done about screen real estate, and that the way people handle documents in the physical world could be used as a source of inspiration. A series of hacks led to the invention of the flip zooming visualization technique, which is the *idea* that is the main contribution of the paper. The idea of flip zooming then resulted in several designs and prototypes, three of which are described in the paper – The Zoom Browser, The Flip Zooming Image Browser, and The Hierarchical Image Browser. (Many more designs and prototypes have been produced based on flip zooming, but they are for the most part outside the scope of this thesis.) But the prototypes were not developed in parallel; the design of each one was influenced by the experience with the previous prototypes, stressing the iterative nature of the design process.

The second paper, *A Framework for Focus+Context Visualization*, built directly on the previous work to come up with some general *princi-*

ples for focus+context visualization. By looking at the many designs and prototypes produced both by ourselves and others in this area, we could find some common factors, and express these in a formal way. Thus the main contribution of this paper is the foundations for a formal system that might eventually be powerful enough to both describe existing focus+context visualization techniques, and produce novel ones.

The third paper, *WEST: A Web Browser for Small Terminals*, took the flip zooming idea all the way back to the hunch stage. The hunch was that flip zooming might not only be useful on desktop computers; it could also be applied to the new breed of handheld computers that were becoming increasingly popular. However, it also soon became evident that flip zooming in itself was not powerful enough to solve the task at hand, in this case presenting a clear view of an ordinary web page on a device with limited processing power and memory, and with a display limited to only 160x160 pixels. But by bringing in a number of other techniques, including proxy pre-processing and text summarization, we could solve the problem, and present the *WEST prototype*, which is the main contribution of this paper. This prototype has attracted enough interest for us to think that it might become a viable commercial product, and *WEST* could become one of the results in this thesis that travels all the way through the spiral and out to a commercial market.

The fourth paper, *Supporting Group Collaboration with Inter-Personal Awareness Devices*, was another return to the hunch stage, but this time without flip zooming as companion. “What would happen”, the hunch went, “if we had a set of devices which would *hum* if they got close to each other”? This turned out to be a not very original hunch (see e.g. [22]), but the subsequent process turned the hunch into something quite different and much more interesting. We produced a hack which proved that this was technically feasible, and eventually came up with a design which we called the *Hummingbird*. Hummingbirds are indeed devices that hum when they get close enough to each other; but the novel element is that we chose to use them as a support for group collaboration rather than an initiator of chance encounters (which has been the goal of most similar commercial systems). Thus Hummingbirds perform a similar task to many desktop-based awareness applications (e.g. [40]) but do so completely outside the scope of WIMP and GUI. This *idea*, of moving awareness information away from the desktop and to the user, is the main contribution of this paper. The Hummingbirds is the second

result of this thesis which has attracted commercial interest, and plans are currently being put into place to turn it into a product.

The fifth paper, *Token-Based Access to Digital Information*, is like the second paper mainly concerned with coming up with *principles*. Through our own implementation of the WebStickers system [24], and through the observation of several other similar systems which allow users to couple digital information with physical objects, we were able to come up with some general results for classifying such systems. The schema of *containers*, *tokens* and *tools*, which aims to capture three significantly different ways in which physical objects can be associated with digital information, is the main contribution of this fifth and final paper.

We can thus see that the results of the five papers in this thesis occupy a variety of different places in Verplank's Spiral. Some stay very close to the hunches and hacks, with the main contribution being an idea arising from this, whereas others are more concerned with general principles. And a few of the results might have the strength to travel all the way out of the spiral and into the marketplace – whether this will actually happen is still too early to tell. In any case, we believe that for scientific work in an artificial science to be successful, it is important to acknowledge all parts of the spiral, although one may choose to focus one's work only at a certain section. The work in this thesis has certainly benefitted from hunches as well as tests, and produced ideas as well as prototypes, designs as well as principles.

But is this work, with its focus on implementation and prototyping, really scientific research, and not “just” product development? The Merriam-Webster's Collegiate Dictionary (online edition) offers several definitions of research, including: “investigation or experimentation aimed at the discovery and interpretation of facts, revision of accepted theories or laws in the light of new facts, or practical application of such new or revised theories or laws”. The work in this thesis has been just such an investigation, with both practical and theoretical components. That some products may indeed come out of the work is true. However, that the nature of the work has been very different from that in a commercial R&D unit should also be clear. It has been very much an exploratory effort, with the intention of testing the limits of the human-computer interface rather than solving any specific problem. Any potential products that have resulted from this have been side effects rather than the main goal. Instead, the main results of the work has been an increased

general understanding of how we can enhance human-computer interaction within or outside the WIMP / GUI paradigms.

5 Conclusion

In this thesis, I have explored a variety of different approaches to make human-computer interaction “break the screen barrier” – to make interacting with computers a richer experience, not limited to the screen of the desktop computer. The work has taken the form of a practical exploration of the possibilities of modern information technology, so that by constructing artifacts and seeing if they work, ideas, prototypes, principles (and perhaps even products) has emerged. This empirical approach to IT research seems to be very much in agreement with the foundations of an artificial science put forth by Bo Dahlbom [14], and also fits well within the scientific agenda of the new informatics [13]. However, it is different from approaches that stress studies and cultivation of current work practice (e.g. [2]), since this work has taken the form of exploratory design of artifacts rather than improvements based in current practice.

Increased understanding of the human-computer interface is today more important than ever before. Spence and Apperley’s vision of the office of the future has still not been realized, although many of the required puzzle pieces have been put in place during the last two decades of research – some of them might perhaps even be found in this thesis. But the use of computers is no longer limited to the office, or even to work. Whereas computers were once specialized and expensive pieces of equipment, computers can now be found almost everywhere. Most families in Sweden have home computers, and electronic entertainment systems are available which approach or surpass the power of current PCs. And this is not all: there are computers everywhere – in cars, dishwashers, TV sets, watches... We are literally surrounded by computers!

Computers can now be made small and inexpensive enough to put inside everyday objects such as furniture and books, jewelry and clothing. Computation has the power to permeate our entire lives, to seep into the very fabric of our existence, just like electric power has managed to do during the last one hundred years. Computers will appear in situa-

tions we never thought possible, aiding, entertaining and comforting us without us even knowing they are there – computers will become invisible. Yet at the same time computers will stay the same: the desktop will not go away. Word processing and calculations will take place on desktop computers until a more powerful alternative presents itself and supplants the desktop, much like computers replaced the typewriter and adding machine that came before them. And there is still much to be explored in the mix between different approaches: on the boundaries between stationary computers and mobile devices, between the digital space and the physical, between powerful number-crunching servers and little things that think. This area, when we have truly “broken the screen barrier”, is where this author believes some of the major innovations in human-computer interaction are still about to happen.

6 References

1. Apple Computer. *Macintosh Human Interface Guidelines*. Reading, MA: Addison-Wesley, 1992.
2. Bergqvist, J. and Dahlberg, P. Scalability Through Cultivation. *Scandinavian Journal of Information Systems*, Vol 11, 2000.
3. Björk, S., Holmquist, L. E., Ljungstrand, P. and Redström, J. Providing Effective Interaction on Small Screens with PowerView. In *Extended Abstracts of ACM SIGCHI Conference of Human Factors in Computing Systems (CHI 2000)*, ACM Press, 2000.
4. Blum, B.I. *Software Engineering: A Holistic View*, Oxford University Press, 1992.
5. Boehm, B.W. *Software Engineering Economics*, Englewood Cliffs, NJ: Prentice-Hall, 1981.
6. Boehm, B.W. A Spiral Model of Software Development and Enhancement. *IEEE Computer*, May, 1988.
7. Bolt, R.A. Put-That-There: Voice and Gesture at the Graphics Interface. In *Proceedings of Computer Graphics Annual Conference Series (SIGGRAPH '80)*, ACM Press, 1980.
8. Brand, S. *The Media Lab: Inventing the Future at M.I.T.* New York, NY: Viking Penguin, 1987.
9. Buxton, W. *Out From Behind the Glass and the Outside-In Squeeze*. Invited speech at ACM SIGCHI Conference of Human Factors in Computing Systems (CHI '97), Atlanta, Georgia, USA, 1997.
10. Card, S.K., Robertson, G.G. and Mackinlay, J.D. The Information Visualizer, an Information Workspace. In *Proceedings of the ACM SIGCHI Conference of Human Factors in Computing Systems (CHI '91)*, pp. 181-186, ACM Press, 1991.
11. Coen, M., (ed.). *Intelligent Environments*. AAAI Spring Symposia Series, Technical Report SS-98-02, AAAI Press, 1998.
12. Cringely, R.X. *Accidental Empires: How the Boys of Silicon Valley Make Their Millions, Battle Foreign Competition, and Still Can't Get a Date*. 2nd ed. HarperBusiness, 1996.
13. Dahlbom, B. The New Informatics. In Ljungberg, F. (ed.), *Informatics in the Next Millennium*, pp. 15-35, Lund: Studentlitteratur, 1999. Originally published in *Scandinavian Journal of Information Systems*, 8(2), pp. 29-48, 1996.
14. Dahlbom, B. The Idea of an Artificial Science. In Dahlbom, B., Beckman, S. and Nilsson, G.B., *Artifacts and Artificial Science*, pp. 2-15, August 1999. Available from: <http://www.informatik.gu.se/~dahlbom/>

15. Davenport, G. Holmquist, L.E., Thomas, M. Fun: A Condition of Creative Research. *IEEE Multimedia*, July-September, 1998.
16. Fertig, S., Freeman, E., and Gelernter, D. Lifestreams: An Alternative to the Desktop Metaphor. In *ACM SIGCHI Conference on Human Factors in Computing Systems Conference Companion (CHI '96)*, pp. 410 - 411, ACM Press, 1996.
17. Fitzmaurize, G.W., Ishii, H. and Buxton, B. Bricks: Laying the Foundations for Graspable User Interfaces. In *Proceedings of the ACM SIGCHI Conference of Human Factors in Computing Systems (CHI '95)*, pp. 442-449, ACM Press, 1995.
18. Gentner, D. and Nielsen, J. The Anti-Mac Interface. *Communications of the ACM*, 39(8), pp. 70 - 82, 1996.
19. Henderson, D.A. and Card, S. Rooms: The Use of Multiple Virtual Workspaces to Reduce Space Contention in a Window-based Graphical User Interface. *ACM Transactions on Graphics*, 5(3), pp. 211 - 243, 1986.
20. Hiltzik, M.A. *Dealers of Lightning: Xerox Parc and the Dawn of the Computer Age*. HarperBusiness, 1999.
21. Ishii, H. and Ullmer, B. Tangible Bits: Towards Seamless Interfaces between People, Bits and Atoms. In *Proceedings of the ACM SIGCHI Conference of Human Factors in Computing Systems (CHI '97)*, pp. 234-241, ACM Press, 1997.
22. Kahney, L. "Hi, Do You Beam Here Often?"; Wired News (Web-based news service), March 25, 2000. URL: <http://www.wired.com/news/technology/0,1282,35090,00.html>
23. Ljungberg, F., Dahlbom, B., Fagrell, H., Bergqvist, M. and Ljungstrand, P. Innovation of New IT Use: Combining Approaches and Perspectives in R&D Projects. In *Proceedings of the Fifth Biennial Participatory Design conference*, pp. 203-210, ACM Press, 1998.
24. Ljungstrand, P., Redström, J. and Holmquist, L.E. WebStickers: Using Physical Tokens to Access, Manage and Share Bookmarks to the Web. In *Proceeding of Designing Augmented Reality Environments (DARE 2000)*, pp. 23-31, ACM Press, 2000.
25. Mackay, W.E., Fayard, A-L., Frobért, L. and Médini, L. Reinventing the Familiar: Exploring an Augmented Reality Control for Air Traffic Control. In *Proceedings of the ACM SIGCHI Conference of Human Factors in Computing Systems (CHI '98)*, pp. 558-565, ACM Press, 1998.
26. Mackinlay, J. D., Robertson, G. G., Card, S. K, The Perspective Wall: Detail and Context Smoothly Integrated. *Proceedings of the ACM SIGCHI Conference of Human Factors in Computing Systems (CHI '91)*, pp. 173-179, ACM Press, 1991.

27. Mann, S. Wearable Computing: A first step toward "Personal Imaging". *IEEE Computer*, Vol.30, No.3, 1997.
28. Muller, M. J., & Kuhn, S., Eds. Participatory design. Special issue of *Communications of the ACM*, 36 (6), 1993.
29. Nixon, P., Lacey, G. and Dobson, S. (eds.) *Proceedings of 1st International Workshop on Managing Interactions in Smart Environments (MANSE '99)*, London: Springer Verlag, 1999.
30. Norman, D. *The Invisible Computer: Why Good Products Can Fail, the Personal Computer Is So Complex, and Information Appliances Are the Solution*. Cambridge, MA: MIT Press, 1998.
31. Rekimoto, J. Time-machine Computing: A Time-centric Approach for the Information Environment. In *Proceedings of the 12th annual ACM symposium on User interface software and technology (UIST '99)*, pp. 45-54, ACM Press, 1999.
32. Rheingold, H. *Virtual Reality*. New York, NY: Simon & Schuster, 1991.
33. Robertson, G., Czerwinski, M., Larson, K., Robbins, D.C., Thiel, D., van Dantzich, M. Data Mountain: Using Spatial Memory for Document Management. In *Proceedings of the 11th annual ACM symposium on User interface software and technology (UIST '98)*, pp. 153-162, ACM Press, 1998.
34. Simon, H. *The Sciences of the Artificial*. Cambridge, MA: MIT Press, 1969.
35. Simon, H. *The Sciences of the Artificial*. 3rd ed. Cambridge, MA: MIT Press, 1996.
36. Sorensen, R. A Comparison of Software Development Methodologies. *CrossTalk*, January, 1995. URL: <http://www.stsc.hill.af.mil/CrossTalk/1995/jan/Comparis.asp>
37. Spence, R. and Apperley, M. *Focus on Information Technology: The Office of the Professional*. Video, Imperial College Television Studio, London, 1981.
38. Spence, R., Apperley, M., Data base navigation: an office environment for the professional. *Behavior and Information Technology*, vol. 1 no. 1, pp. 43-54, 1982.
39. Spence, R. and Apperley, M. *The Bi-focal Display*. Video, Imperial College Television Studio, London, 1983.
40. Tollmar, K., Sandor, O and Shömer, A. Supporting Social Awareness @Work, Design and Experience. In *Proceedings of CSCW 96*, pp. 298-307, ACM Press, New York, 1996.
41. Weilenmann, A. and Holmquist, L.E. Hummingbirds Go Skiing: Using Wearable Computers to Support Social Interaction. In *Pro-*

- ceedings of Third IEEE International Symposium on Wearable Computing (ISWC '99)*, pp. 191-192, IEEE Press, 1999.
42. Weiser, M. The Computer for the 21st Century. *Scientific American*, 265(3), pp. 94-104, 1991.
 43. Wellner, P., Mackay, W. and Gold, R. (Eds.) Computer-Augmented Environments: Back to the Real World. Special issue of *Communications of the ACM*, 36(7), 1993.