Evaluating Unattended Technology, a Subset of Calm Technology

by

Vladimir Levin

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Mathematics
in
Computer Science

Waterloo, Ontario, Canada, 2008

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Information is a central theme of the twenty-first century. This is evident in the fact that everyday objects are being augmented to provide information. Thus, ubiquitous computing – providing information using everyday objects – becomes increasingly popular.

The problem is that information requires attention for acquisition. Hence, ubiquitous computing puts a strain on attention, which is limited. There are many innovations that attempt to solve this problem; this thesis focusses on one: calm technology, which was introduced to interface design by Mark Weiser. Calm technology attempts to reduce the attention required to acquire information. Ideally, calm technology would provide information without requiring any attention. I call this technology unattended.

Calm technology research, however, typically provides little evidence showing that calm artifacts reduce the amount of attention required. Moreover, evaluations that are conducted on individual artifacts often fail to generalize. That is, evaluations only apply to the artifact that is evaluated. They do not identify properties of the artifact that make it calm.

In this thesis, I design and conduct a dual task experiment. The results of the experiment indicate that users can perform an attention saturating primary task, and acquire information from a calm artifact not involved in the task, without sacrificing performance on the primary task. Thus, the artifact does not require any attention, as can be measured by the experiment, while providing information. Thus, the artifact is unattended, which provides an existence proof for unattended technology.

Acknowledgements

I would like to thank my supervisor, William Cowan, for sharing his wealth of knowledge, and being patient, with me.

Additionally, I would like to thank John Beatty and Edward Lank for reading and correcting numerous errors in my thesis.

Also, I would like to thank all of people who participated in my experiments for their time and attention.

A special thanks goes to the juice fund guy for providing much needed, and inexpensive, nourishment in the lab.

Most of all, I would like to thank Vedrana Toperic for being patient with and supportive of me as I worked on my experiments and my thesis.

Dedication

To V.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Many call the twenty-first century the century of information, whatever that means. One thing is certain, however. Things are increasingly evaluated in terms of their information content: images are the information they represent; the genome is the information it contains; employees are the information they know.

Thus, the objects in our environment are continually updated by adding information to them. Ovens, for example, are augmented with programs that cook automatically, with the side effect of showing users the time and date. Refrigerators are proposed that know the caloric content of food they contain, and that remind users when, for example, fresh milk is needed. Such rudimentary capabilities are, of course, the tip of the iceberg. As computational objects are made to communicate they offer their users many of the information capabilities that were previously offered by a host of servants and services.

The end-point of this evolution is called pervasive or ubiquitous computing, which occurs when most objects in the environment have internal computation, provide rich computational output as part of their presentation, and accept input with deep computational consequences. The result is a severe challenge to the user's capabilities. Only a few people work in such environments: traders in financial institutions, or air traffic controllers, for example. The skills they require are exceptional and burn-out is common.

Herbert Simon long ago explained why: information consumes the attention of its recipients, which is strictly limited, "... a wealth of information creates a poverty of attention." What underlies this reality is simple. The output of computation is information, which requires a limited resource, sometimes called attention, in order to be acquired by the user. One result of using this limited resource is the transfer of information to long-term memory. Only information successfully incorporated in long-term memory is easily combined with perception to produce decisions and plans of actions.

Understood thus, an overabundance of information puts unsustainable demands on limited resources. (Scholars understand this problem well. Time spent reading research papers, of which there is never enough, merely transfers information from external sources to long-term memory for later use.) Widening this channel is thus an essential prerequisite for effective exploitation of ubiquitous computing. Exploration of this problem is the core of this thesis.

An obvious place to seek a solution is information visualization, the goal of which is visual encoding of information so as to ease access. Successful visual encoding, presumably, minimizes the attention consumed in acquiring information. This is one of the ways that more information can be provided for the same amount of attention.

Visualization supplies the user with information in return for his attention. Indeed, exploring information using visualization is a totally engrossing task, which does not admit distraction. Its benefit is not reducing attention, but giving more information for the same expenditure of attention. The converse, giving the same information for a smaller expenditure of attention – giving information, for example, to a user occupied with an unrelated task – is another approach for increasing attention supplied to the user. In the limit, this approach would amount to providing information to a user whose attention is fully occupied with another task. Such a goal is the focus of this thesis.

Calm technology, which has always existed, and which was introduced to the user interface by Mark Weiser, comes close to sharing this goal. At present, there exist about a dozen studies that call themselves examples of calm technology. Each introduces one or more calm artifacts, each one a user interface that is said to embody one or another aspect of calm technology. While most of these studies are interesting in themselves, theoretical abstractions linking them into examples of a general domain are, more or less, absent.

These studies are discussed in detail in chapter 2. Each includes the claim that the interface developed is 'calm'. These claims are hard to evaluate, because theoretical research justifying a definition of 'calm' seems not to exist. Implicit in most research, however, are two components of 'calm'. First, calm artifacts are aesthetic, usually expressed negatively as meaning that they do not distract or annoy. Second, they provide information without forcibly claiming attention.

A modest interpretation of the second point is that the user's attention is managed by the user, not by the interface. In other words, the user controls whether or not to allocate attention to acquiring information. Thus, visualization differs from calm technology just as a movie differs from an art gallery. In a movie, what is on the screen changes constantly. Misdirected attention can easily miss critical information. For this reason talking is discouraged and the room is dark, to minimize diversion of attention. The viewer, in essence, surrenders attentional control to the movie's director. In contrast, in an art gallery, the user freely chooses where to direct his attention. A calm artifact is similar. It does not force the user to attend to it. The decision to allocate attention is left with the user.

Clearly, attention is an important concept in calm technology. But, what *is* attention? The Shorter Oxford Dictionary defines it using 'attend', which has seven non-obsolete definitions, one of which is circular with attention. The commonality, when attention is discussed in the context of calm technology, is that a resource, a memory buffer or a processing channel, is limited, and attending means allocating some or all of it to a task. There is a set of resources, and all of them are more or less limited. When attention is paid to an object, some or all of the resources are allocated to that object, and, as a result, these resources are not available for other tasks. In research on calm technology, it is understood that calm artifacts require only a little attention to provide information. In the extreme case, a calm artifact would not require *any* attention. That is, the limited resources are not used when acquiring information from the artifact. In this thesis, such artifacts are called unattended artifacts. Unattended technology, then, refers to an extreme case of calm technology, where attention is not required to interact with it.

Fortunately, a well-accepted methodology for deciding whether an artifact requires attention exists: the dual task experiment. The user's attention is saturated by one task when a second task is presented. If the second task requires attention, the user has only two choices. He can continue to perform the first task, ignoring the second. Or, he can perform the second task, sacrificing performance on the first. Such an experiment is described in chapter 3, based on a typical calm artifact. It provides a concrete method for evaluating unattended artifacts.

The experiment is then used to evaluate a calm artifact, which is designed to require no attention. The results of the experiment, which show that non-trivial information can be acquired and processed without attention, are discussed in chapter 4.

The experiment is a proof of existence of unattended technology, and, hence, calm technology, which makes it possible to treat calm technology as a science. That is, questions about calm technology, and unattended technology, can be framed in terms that admit an unambiguous answer: questions such as "How much information can be provided using unattended artifacts?" and "Can precise information be provided using unattended artifacts?" Having a proof of existence provides some interesting research opportunities, which are discussed in chapter 5.

Chapter 2

Previous Research

This chapter discusses previous research in calm technology. It begins with a history of calm technology. Then, it introduces a definition of calm technology, which comprises a set of criteria that is common to most calm technology research. After discussing the definition, it describes papers that introduce and evaluate calm artifacts. Of particular interest are evaluation methods that show how successful a calm artifact is with respect to the definition.

2.1 History of Calm Technology

Calm technology has always existed. Most of the time, it is easy to recognize a calm artifact when it is seen, but it hard to pick out properties that make a general artifact calm.

Examples of calm technology are numerous. For instance, when one walks outside, he gathers information about the weather. This is evident, because once inside and asked about the weather, one can answer accurately, even if attention was not devoted to inspecting the weather. The information is acquired simply by being outside. This means that the weather information is calm. Even when working in an office, one can acquire weather information from a window. This makes a window an example of calm technology.

Calm technology does not have to be visual. For instance, when driving a car, the driver rarely pays attention to the sound of the engine. However, if something is wrong with the engine and the sound changes, it is noticed immediately. In some way the driver always heard the sound, which required no response. But when it changed, the information changed, and the driver received important information: something was wrong with the engine.

Calm technology also long existed in computing technology. One of the best examples is the input to a typical computer, its keyboard. The keys of a typical keyboard have a certain feel, which was designed in. That is, when a key is pressed, a quiet clicking sound occurs and the finger pressing it feels a pattern of resistive force. This informs the fingers of the person typing that the key has been pressed successfully. When the user does not receive correct feedback from the key, he stops typing and tries to identify the problem. This requires attention. It is evident that attention is required only when the key is pressed incorrectly. Hence, the information provided when the key is pressed correctly must be calm, which makes a keyboard an example of calm technology.

It is important the keys do not produce a loud bang every time they are pressed, because that would be distracting. The sound keys make is loud enough to hear, but not loud enough to disturb the user.

It is clear that calm technology can be found everywhere. Most of the time, an artifact, which is designed to be pleasant to interact with and not to be disturbing, is calm.

Research designed to illustrate aspects of calm technology, was done by Mark Weiser, who gave calm technology its name, after identifying some of its properties in everyday objects [24]. He then introduced LiveWire (Figure 2.1), the canonical example of calm technology. LiveWire incorporates several aspects of calm technology that Weiser considered important. In particular, Weiser considered the important properties of calm technology to be aesthetic value and ability to provide information while remaining in the “periphery of attention”, meaning while

Figure 2.1: LiveWire: canonical example of calm technology. The camera is positioned near the floor and is facing up. The 'caution' sign is on the wall, and LiveWire is located on the ceiling.

the user is not compelled to pay attention to an artifact.

Here is how LiveWire works. A network cable is connected to a network interface card (NIC). The CPU on the NIC is connected to a stepping motor, which has an 8 foot plastic string dangling from it. The CPU gathers statistics on the network traffic passing through the cable. From them it calculates the density of traffic, and rotates the motor by an amount depending on the density. The more traffic on the network, the faster the motor rotates, and the faster the wire spins.

The result is a device that creates an animated visual representation of network traffic. When the network traffic is low, the string is motionless. However, as network traffic increases, the string starts spinning faster, making visual motion and a whirring sound.

According to Weiser, LiveWire is primarily an aesthetic object, a work of art, which secondarily allows the user to know network traffic, while expending minimal effort. It assists the user by augmenting an office with information about network traffic. Essentially, it moves traffic information from a computer screen to the 'real world', where the user can acquire information from it without looking directly at it.

Calm technology can be separated into two categories. LiveWire is an example of a trend in calm technology research, which introduces artifacts that provide information in an aesthetic form. An office window is also an example of this. It is aesthetic and it provides information while requiring little attention. These artifacts are also known as ambient displays, peripheral displays, or informative art, which is the name used in this thesis. The papers that fall into this category have a common goal: to create a work of art that provides information to the user, while requiring little attention from the user.

A second category in calm technology, which includes, among others, things like key presses on a keyboard, is developing artifacts for augmented reality, a sub-field of computer graphics that can benefit significantly from calm technology. Typically, calm artifacts that are developed for augmented reality provide contextual information that assists the user with a task at hand. They must be unobtrusive, and the user must be able to acquire information from them with minimal effort.

These categories encompass most of calm technology research. Research papers in each of these categories are described below. By discussing the papers in each of the categories, common elements of calm technology will become evident. In turn, this will lead to a precise definition of calm technology.

2.1.1 Informative Art

Informative art is one of the large categories of calm technology. The research in informative art concentrates on introducing artifacts that are visually pleasing, and provide information while remaining in the periphery of the user's vision. There are many papers that introduce and describe examples of informative art [3, 4, 7, 9, 10, 11, 19, 22, 23, 24]. Although each of the papers introduces a unique artifact, with its own properties, each artifact contains a set of similar characteristics. These are evident in the papers I describe here, which are examples of typical informative art.

Holmquist and Skog described several information displays, collectively called Informative Art [9] (not to be confused with informative art as a category of calm technology). These displays depict information such as time, weather conditions, and earthquake history in certain regions in a form of abstract paintings. The authors picked several art styles, styles adapted from artists such as Piet Mondriaan, Bridget Riley, and Andy Warhol, and encoded information into them to produce displays that are aesthetic and provide information. These displays are meant to be installed in a home environment, just as the art they imitate might be. In addition to arguing that the displays are aesthetic, the authors claim, indirectly, that the displays require little attention, by stating that they had to redesign several of the displays (reduce brightness, or speed of animation), because they required too much attention.

Cadiz et al. introduced Sideshow, an information bar that is located on the side of the user's desktop [3]. It displays text and images and is able to provide a variety of information, from time and weather to online status of other users. This information was provided using text and images. For instance, "Sarah is online" is displayed on Sideshow, next to an image of a full moon. This informs the user that another user, Sarah, is currently by her computer, and that the moon is the full phase. Although the information is located on the computer screen, not in the 'real world', the authors consider this a calm artifact. That is, according to the authors, it is present in the periphery of the user's vision, and it requires little attention to be useful. In the author's words, the users are 'aware' of the information Sideshow provides. Note that although the paper does not discuss whether Sideshow is aesthetic, it is clear from the discussion of its design that the goal was to make it as unobtrusive as possible.

Tobias Skog introduced Activity Wallpaper, a display that contains a visual representation of ambient noise within a cafe over a period of a week [19]. It is visualized as stylized bar graph. There are seven bars, one for each day of the week. The bar height represents the amount of ambient noise during the day. Finally, each bar is rendered as a collection of small aligned circles. The more noise the cafe has, the higher the bar extends. The authors claim that Activity Wallpaper acts like a decoration, a mural, in a public setting. By providing historical information about the place, it "gives users a feel for how the place was inhabited" [19].

These examples demonstrate typical informative art research. The papers concentrate on developing artifacts that are, to some degree, aesthetic, provide information, and require little attention from the user. These are the common elements in most informative art papers. By combining these elements into an artifact, informative art is made.

2.1.2 Augmented Reality

Augmented reality is a well studied field of computer science. Its goal is to enhance the user's environment by providing information that would otherwise only be available on a computer screen. This is sometimes accomplished using translucent glasses that overlay rendered images onto the environment seen through the glasses. In other instances, environments are augmented by introducing artifacts that provide information.

Calm technology can be useful in the latter method, because it has a similar goal: to provide information to the user in an aesthetic way, while requiring little attention. In fact, there are many papers that do introduce calm artifacts to augment the user's environment by providing information that can assist the user with the task at hand [5, 10, 11, 13, 17, 23, 24]. Note that some of these papers overlap with informative art, because some artifacts include elements of both informative art and augmented reality. Below, I describe two papers that are typical examples of this type of calm technology.

MacIntyre et al. introduced an augmented office display [13]. They used a projector and a large screen to display a montage of the current work environment. The screen was located behind the user's computer monitor and displayed information such as history of commands and data related to the current task. The display brings non-trivial information from the computer screen into the real world. Although the screen is just a large computer monitor, it can provide information while being in the user's periphery of vision. The authors claim that the users are aware of the information. They also provide a way to montage the work history in a way that users understand quickly. In essence, the user does not have to devote his full attention to the large screen in order to acquire information from it. In turn, this means that the user's environment is augmented by the display, which requires a small amount of attention to be useful.

Vogel and Balakrishnan introduced a public display, called the "hello wall", that had the capacity to display both public and private information [23]. The screen monitored the user's proximity to the screen and displayed different information depending on the user's location. That is, while the user was far from the screen, it displayed information such as weather conditions and the volume of people in an office. Similar to informative art, the information presented in the "hello wall" was a visualization in forms of simple graphs and abstract art styles. However, when the user approached closer, it displayed more private information, such as the user's calendar and a private messaging system. In general, when the user was close to the "hello wall", it turned into a regular computer monitor. Hello wall was then an calm artifact when the user was far away from it, but became a regular computer monitor as the user approached. This augmented the user's environment by providing varying information in an aesthetic way depending on the user's location.

As these two projects show, calm technology in augmented reality is characterized by providing contextual information while requiring little attention. Indeed, augmented reality is a research field that can benefit significantly from calm technology research.

Clearly calm technology is useful, but what makes a regular technology calm?

To answer this, we need to analyze artifacts that exists in calm technology research.

2.2 Definition of Calm Technology

From the papers introduced in the previous section, an understanding of what makes an artifact calm begins to emerge. There are several criteria that are common to all calm technology papers.

1. It is aesthetic. This is mostly interpreted to mean that the artifact is visually pleasing, particularly in informative art papers. However, not all papers make this a requirement. Most augmented reality papers that develop calm artifacts only state that the artifact has to be unobtrusive. What is evident, and what I interpret this to mean, is that the artifact is aesthetic when it fits into the environment without distracting the user. It can exist in a visual environment without drawing unnecessary attention because of its appearance.

2. It provides information. Calm artifacts must provide information, which is their purpose. All calm technology papers that introduce artifacts encode some type of information into them.

3. It requires little attention for information acquisition. This criteria separates calm artifacts from ordinary visualization, which requires full attention to be useful. Calm artifacts can provide information to a user who is occupied with another task. This is explicitly stated in some papers, and implied in others. However, the fact that the attention required for a calm artifact to be useful is small consistently appears in calm technology research.

When designing calm artifacts each of the criteria needs to be considered. If one criterion is left out, the artifact will no longer be calm. This is because the criteria defined are common to all calm artifacts.

First, the aesthetic criterion has to be considered. However, because this criterion is subjective, the artifact simply needs to be unobtrusive enough for users to accept it in their environment. Although it can be visually pleasing, as are most examples of informative art, it can also simply fit into the environment, as is typically done in augmented reality.

Second, the information content of an artifact has to be considered. Although the amount of information can be measured fairly easily, there is no set minimum amount of information that a calm artifact has to provide. When designing a calm artifact, it is enough to include one piece of information to satisfy the criterion.

Lastly, the amount of attention the artifact requires needs to be considered. Clearly the amount of attention a calm artifact requires in exchange for information is important. Estimating the attention requirement of an artifact reveals whether it is an instance of calm technology or an ordinary visualization. To measure the amount of attention an artifact requires an evaluation is needed. Below, formal evaluation techniques that have been used in calm technology research to determine whether an artifact is calm are explored.

2.3 Evaluation of Calm Technology

There are only a few research papers that evaluate calm artifacts. The evaluation methods differ in technique and the type of properties that they discover. The methods, with their advantages and drawbacks, are listed in this section.

2.3.1 Designer and User Evaluations

One of the themes in calm technology is methodologically designing calm technology to fit calm technology's criteria, as authors define it, which generally includes calm technology's criteria as I define it. For instance, using user interviews and questionnaires can reveal that the user does not like an artifact, because it forces him to pay attention to it. As a result, such artifacts cannot be calm: they require too much attention. Hence user interviews, and questionnaires, can identify artifacts that are not calm.

There are numerous papers that take this approach to developing calm technology [3, 4, 10, 16, 20]. Two papers that are representative of the theme are described here.

Iqbal et al. introduced methods that integrate disciplines other than computer science and information visualization in order to design successful calm artifacts [10]. They have used concepts from human-computer interaction, cognitive science, anthropology, and social psychology to design questionnaires and interviews that assist them in designing calm artifacts for user collaboration. Although the authors did not elaborate on the details of questionnaires, they did say that by conducting evaluations with real users of the system, they were able to guide the design of their system to fit the criteria of calm technology.

Stromberg et al. turned to role-playing as a tool to design calm artifacts [20]. By letting users play out interactive scenarios and improvise with mock calm artifacts (some artifacts were represented by pieces of paper, others were represented by other actors), they have discovered how users would interact with real calm artifacts. The paper addresses the early design stage of calm technology. The authors found that role-playing is useful in determining exactly what the user expects from a calm artifact and how he interacts with it. In turn, this assists in the design of calm artifacts.

These methods are partially effective. They ensure that the user will be satisfied with a calm artifact, because the user is involved in the design process. However, these methods are typically applied in the design stage of calm artifacts. That is, once an artifact is implemented, researchers use interviews and questionnaires to evaluate long-term installations of calm technology.

An exception to this is the work of Mankoff et al. They developed a heuristic method for detecting usability issues in calm artifacts at any stage of development [16]. The success of Mankoff's heuristic was confirmed by another researcher, who claims that the heuristic is a good way to determine whether an artifact is calm [4]. The heuristic is described in more detail below.

Mankoff's heuristic

Mankoff's heuristic contains twelve categories. The three categories that correspond to calm technology's criteria are "aesthetic and pleasing design", "sufficient information design", and "'peripherality' of display".

The "aesthetic and pleasing design" category states that the display should be pleasing when placed in its intended environment. This is a stronger criterion than my aesthetic criterion of calm technology, because it requires the display to be visually pleasing. Hence, it can certainly determine if the display does not fit into its environment.

The "sufficient information design" category states that the display should provide just enough information. Too much information would appear too crowded, and too little information would make the display less useful. This corresponds to my second criterion of calm technology: the fact that it provides information. However, in Mankoff's heuristic, the condition is more restrictive. That is, there is both an upper and a lower bound on the amount of information the display should provide. However, the bounds are not precise, and are unnecessary. For instance, the feedback a key provides when pressed amounts to one bit of information: yes, the key was pressed. However, this information is just as useful as any other piece of information because it allows the user to type confidently and without interruptions. If an artifact contains too much information, on the other hand, then it would not be possible to acquire it while expending little attention. That is why my second criterion of calm technology is simply to provide information.

Finally, the "'peripherality' of display" category states that the display should be unobtrusive and remain so until the user pays attention to it. This category roughly corresponds to the third criteria of calm technology. That is, while the category states that the display should not draw attention to itself, it does not say that the display should provide information while it is in the periphery of the user's vision. Most calm technology research does imply that information should be provided while the user expends little attention. That is why the criterion is included in my calm technology definition.

The other categories in Mankoff's heuristic do not appear in calm technology's criteria, or in most other calm technology papers. For instance, the "visibility of change" category says that the artifact's transition from one state to another should be perceptible by the user. However, this seems to contradict the 'peripherality' of display category. If the changes on the display are noticeable, they might draw the user's attention to the display. Another category states that the information encoding in an artifact should be intuitive. This is implied by my third criterion of calm technology: if the user can acquire information from a display while expending little attention, then the encoding is intuitive. On the whole, there are only three categories that correspond to calm technology's criteria. All others are either vague or implied by one of my criteria.

Overall, according to the authors and as confirmed by another researcher, Mankoff's heuristic method allows 3-5 evaluators to detect 40-60% of usability issues, which makes it a successful method for detecting artifacts that are not calm.

On the whole, designer and user evaluations are an effective tool for determining whether an artifact is calm. These evaluations are good at pinpointing problems with an artifact, so that the designer can fix them. However, they do not specifically address attentional demand on a calm artifact.

2.3.2 Long-Term Evaluation

Another common method for evaluating a calm artifact is to place it into a real environment for a certain duration, usually more than a week, and gather users' reactions to the artifact at the end of the study, as was done in [3, 4, 11, 22]. By performing a long-term study, researchers are able to determine whether the tested calm artifacts were successful, and how they were used, in the real world. Here, I describe three papers that use a technique common to calm technology papers that perform a long-term evaluation.

Consolvo and Towle introduced CareNet, their calm artifact for the home [4]. CareNet informs care givers of medication that an elderly person needs by embedding symbols representing this information in a picture frame containing the elderly person's photograph. The researchers initially applied Mankoff's heuristic to find usability issues in their system. They have confirmed that Mankoff's heuristic can find a large set of usability issues by applying it to CareNet. That is, the heuristic revealed usability issues in CareNet, which the researchers fixed before doing more evaluation.

The researchers also performed a three week study by placing CareNet into an actual home. By performing interviews with users after the study, they found a further usability issue not detected by Mankoff's heuristic. In low light conditions, the display was too bright, which was distracting. This is an important issue, but it was not detected by Mankoff's heuristic.

Van der Hoog et al. introduced and evaluated Gustbowl [22]. Gustbowl was designed to provide basic communication between mother and son. It consists of two bowls, one placed in the mother's house, and one in the son's. When an object is placed into one of the bowls, the other bowl displays a picture of that object and begins to wobble. This provides a "Mom, I'm home!" type of message.

However, while performing a long-term evaluation, the researchers found that users did not use Gustbowl solely for its designed purpose. Gustbowl is not meant to be noticed. It is meant to provide subtle cues that a person arrived in his house. It is meant to contain objects like car keys. However, one of the mothers started collecting small trinkets that reminded her of her son, and placing them into the bowl for the son to see. The son 'replied' by placing objects he collected during his day into the bowl. This provided more than "I'm home" type of communication. According to the mother and son, they felt much closer to each other.

Jafarinaimi et al. developed Breakaway to change human behaviour [11]. Breakaway is a small figurine, which represents a sitting human shape, placed on the user's desk. The longer the user works at her desk, the more the figurine slouches, informing the user to take a break.

By performing a long-term study, the researchers found that the user's behaviour changed for the better: breaks were taken more often, which supported a healthy lifestyle.

Evidently, long-term studies can reveal a variety of problems and uses of calm technology. They can determine actual uses of a calm artifact, whether real users use the artifact, and whether calm artifacts perform as intended. Hence, long-term evaluation is an essential part of designing practical calm artifacts.

However, long-term studies have flaws. First, they are long and, ipso facto, expensive. Researchers have to be confident that the calm artifact is successful before they are willing to perform a long-term study. Second, long-term evaluations generally do not answer questions about an artifact's attention requirement or comprehension. They simply prove that one particular artifact is successful: the user likes, or uses, the artifact. They do not help in future design of new artifacts, because they do not provide insight into particular elements of calm artifacts that contribute to their success. Also, they do not address the question of attentional demand. Hence, they cannot determine whether a calm artifact is unattended.

2.3.3 Experimental Evaluation

Experimental testing is a third evaluation technique in calm technology research. It allows precise measurement of both the attention requirement of a calm artifact and the user's comprehension of information provided by the artifact. Researchers can vary a single parameter of the artifact, and measure how it affects comprehension and attention requirement. This has been done for some calm artifacts [5, 8, 15]. I provide two examples that capture the essence of experimental evaluation of calm technology.

Costanza et al. developed eye-q, a display mounted in the user's glasses that notifies the user about predefined events. Eye-q can notify the user about events, such as incoming e-mail or calendar notifications. This is done by displaying symbols and patterns on a glasses-mounted LED display.

The researchers investigated the effect of LED brightness, as well as speed of changing patterns, on detection and attention by conducting an experiment. They found that bright LED settings caused the user to detect patterns better, but they saw a decrease in performance on the primary task, meaning the attention requirement also increased. Dim LED settings with slow changing patterns required less attention, but the user was slower at detecting that a pattern existed.

By performing experimental evaluation, researchers were able to measure the effect of a small set of settings, brightness and speed, on detection and attention requirement of a particular calm artifact.

The example of an experimental approach that serves as the main motivation for this thesis is the work of Maglio and Campbell [15]. Their idea was to have the user perform an attention saturating task, text editing, while a calm artifact, scrolling text, appeared. Using this approach, the researchers measured two things:

1. The attention requirement of the scrolling text. This was measured as an amount of degradation in text editing performance from the base case, where the scrolling text was absent.

2. The comprehension of the information the text provided. This was measured as a performance on a questionnaire that followed the experiment.

This approach specifically addressed the question of attention requirement and comprehension of a particular calm artifact. The researchers investigated the effect of text motion (with and without pauses) on performance and found that text motion without pauses required less attention, while providing the same amount of information. This allowed the researchers to objectively evaluate a calm artifact.

In general, experimental evaluation of calm technology allows the researchers to identify specific features of an artifact that contribute to its success. Furthermore, experiments can investigate how a change in one parameter affects the user's attentional demand and comprehension of the artifact. This indicates that experimental evaluation is a good tool to use in order to develop a concrete understanding of calm technology. In particular, experimental evaluation can be applied to calm artifacts to determine whether they are unattended. If a calm artifact does not require any attention, as measured by the experiment, to provide information, it is unattended.

2.4 Evaluation of Unattended Technology

The limiting case of calm technology, unattended technology, does not require any attention to provide information. This is an interesting technology to investigate. However, looking at the examples of calm artifacts discussed above, it is not evident whether unattended technology is possible. It seems to exist – for instance, a key press does not require any attention to provide feedback to the user – but experimental evidence is missing.

Attention demand is the most important factor to investigate in unattended technology, because it is what separates it from calm technology. Because experimental evaluation is the only technique used in calm technology that can detect attention claims accurately, I use it as the point of departure for showing that unattended technology is possible.

The question that needs to be answered if unattended technology is going to be investigated further is whether it is possible to have unattended technology. This is the question this thesis investigates.

In the following chapters, I design and conduct a dual task experiment, using an approach similar to Maglio and Campbell. The results of the experiment reveal that users can acquire information from an artifact without using any attention. In turn, this shows that the artifact is unattended, which provides an existence proof of unattended technology.

Chapter 3

Experiment Design

The remainder of the thesis concentrates on showing that unattended technology, a limiting case of calm technology, is possible. This is done using a dual task experiment.

The chapter begins with an overview of a model of the passage of information from sensation to long-term storage. After discussing the steps involved in the process, it talks about the pilot study that was conducted, including its results and implications. Then the methodology for the experiment used in this study is discussed. Finally, implementation details of the experiment are given.

3.1 Memory

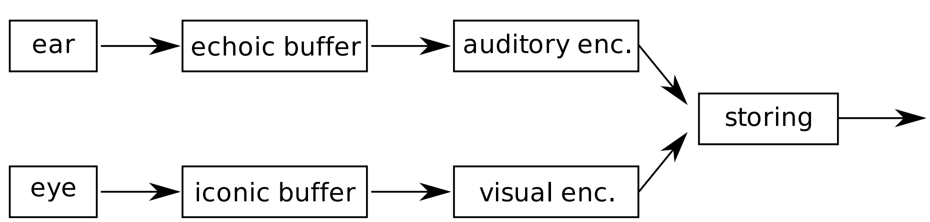

Figure 3.1: Memory process

Long-term memory allows storage and retrieval of information. Any information that the user acquires, and uses at a later time, is stored in memory. The steps required to form a memory are shown in Figure 3.1.

Initially, information is acquired from the environment using sensors, such as eyes and ears. This uninterpreted information is stored in a sensory buffer, which retains the information for a short time, generally less than one second. Sensory buffers for visual information are called iconic buffers. For auditory information, sensory buffers are called echoic buffers.

Information in the sensory buffer is encoded, using either visual or auditory encoding, depending on the sensor that gathered the information. Opinions vary with respect to how much encoding uses resources common to different modalities. In my model, the encoding mechanisms are independent, as suggested by [2] and as observed in my experiments.

For instance, when one hears a sound, it is separated into different streams based on the source of the sound. If one of the streams contains speech, it is then separated into words or phrases, which are, in turn, processed for meaning. The words and phrases may be then processed at a higher level to comprehend the meaning of the sentence.

A similar process happens for visual encoding. The image is segmented into objects based on visible boundaries and object distances. Each individual object can then be processed further in terms of shape, colour, and location in three dimensions.

The encoded information is stored in memory for later retrieval. It can also be further enhanced with context information, but the steps described here are sufficient for the thesis.

Note that there is a set of limited resources that is typically used during the encoding step in this processes [1]. But, the resources allocated for each type of encoding are not necessarily the same. For instance, both visual and auditory encoding typically require some limited resources. Some of the resources are the same for both modalities, and some are not. Consider looking at a picture and listening to music at the same time. This can be done easily, because the limited resources used for each of the activities are different. However, reading a text and listening to someone speak is much harder. That is because the resources needed for language comprehension is used for both tasks. Since it is limited, it is hard to perform both tasks well.

3.2 Pilot Study

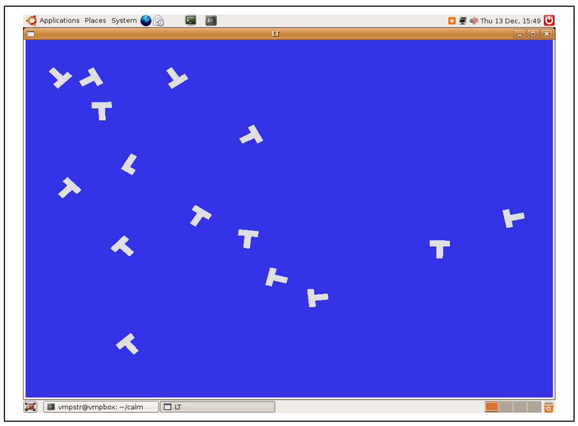

Figure 3.2: Ls and Ts: the first experimental attempt.

The pilot study that I have conducted was a visual search experiment, which is described in detail below. The results of the pilot study show that it is possible for information to be placed into the iconic buffer without requiring attention.

3.2.1 Design

The task done by the user was a search task. The user had to determine whether or not an L was present in a field of Ts, doing so as quickly as possible. The letters were rotated and placed randomly on the screen to make the task attention saturating, in the sense that concurrent tasks increase search time.

The number of letters varied from 5 to 25 in increments of 5, because the important dependent variable is search time per distractor.

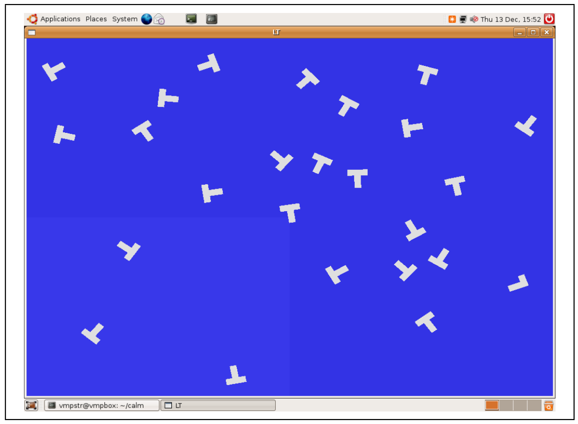

Figure 3.3: Ls and Ts: bottom-left quarter is highlighted.

A calm artifact was added to the experiment in the form of a small gradual brightening of the quarter of the screen in which the target (L) would appear on the succeeding trial, as shown in Figure 3.3.

Because it was necessary to measure both the amount of attention needed to perceive the highlighting, and the amount of help it provided, the highlighting provided information about the succeeding trial. That is, if the bottom-left quarter is currently highlighted, then the target can still be anywhere on the screen. However, after the user selects a choice (whether the target is present or not), the next screen can only contain the target in the bottom-left quarter.

This yielded four distinct cases.

1. Base case. A base case trial had no highlighting, and followed a trial with no highlighting.

2. Help case. A help case trial had no highlighting, and followed a trial with highlighting. In this case, the user could have acquired knowledge from previous trial, allowing him to search only one quarter of the display.

3. Distraction case. A distraction case trial had highlighting, and followed a trial with no highlighting. In such trials, the user had to search the whole screen for the target, but would only have to search a quarter of the screen

Type	Mean	Standard Error
Present case const.	657 ms	123.1
Present case slope	61 ms	7.6
Absent case const.	928 ms	137.4
Absent case slope	121 ms	8.3

Table 3.1: Base case times

on the next trial. This case measures how much extra time the user requires to encode and store where the highlighting is located.

4. Mixed case. A mixed case trials combined elements of the help and distraction cases. There was highlighting present, and it followed a trial with highlighting present. The user must search only a quarter of the current screen, and must encode and store the highlighting.

Each case was further split into an absent case and a present case. The absent case did not have the target present, whereas the present case did.

3.2.2 Results

Both the slope and the constant of the regression line are important for interpreting the results. The slope measures how much time the user spent on each distractor while searching for the target. The constant measures the per trial overhead time.

The results of the base case are listed in Table 3.1. Note that the present case slope is 61 ms, which is an appropriate time for this type of search. Also, the absent case slope is roughly twice the present case slope. This is also appropriate. That is, on average, the user had to search half as many distractors in the present case, because he could have found the target at any time during the search. These two items indicate that the search task was done sequentially, which was the goal.

The results of other cases are listed in Table 3.2, and are discussed in terms of their difference from the base case times.

Type	Help Ave.	Help S.E.	Distraction Ave.	Distraction S.E.	Mixed Ave.	Mixed S.E.
Present case const.	630 ms	113.5	759 ms	133.0	498 ms	141.7
Present case slope	36 ms	6.5	61 ms	7.3	61 ms	7.3
Absent case const.	903 ms	158.7	956 ms	112.1	820 ms	135.5
Absent case slope	85 ms	9.4	123 ms	7.5	110 ms	9.3

Table 3.2: Search task times

First, the help case results show that the highlighting assisted the user. This means it successfully provided information. The important measure for the help case was the slope, because if highlighting helped, then the user would only search a subset of distractors. Since the number of distractors itself did not change, the average time per distractor would decrease.

The difference in the slopes for the help case was -25.16 ms (s.e. 7.54, $t = -3.34$, $p = 0.001$) in the present case, and -35.35 ms (s.e. 6.77, $t = -5.22$, $p < 0.0001$) in the absent case. This is a statistically significant difference, which indicates that the highlighting helped the user perform the search task. The help was, however, not perfect, in the sense that the user searched more than a quarter of the screen. This could be due to the user not being sure about the boundaries of the highlighting, or due to the fact some of the time the user did not use the highlighting and some of the time he did.

Second, the user did not spend extra time to remember the position of the highlighting, as indicated by the distraction case results. In the distraction case, the important measure is the constant term, not the slope, because it would take a constant attention, and time, per trial to remember where the highlighting was. The number of distractors, or time spent per distractor, is not affected by remembering the highlighting.

The difference in the constant terms for the distraction case was 101.81 ms (s.e. 151.8, $t = 0.671$, $p = 0.503$) in the present case, and 28.58 ms (s.e. 116.6, $t = 0.245$, $p = 0.807$) in the absent case. Neither difference is significant. In particular, the 95% confidence interval for the difference in both the present and the absent case contain 0, which means that the difference is not significantly different from 0. However, judging from the standard error of the results, the experiment could have only detected times above about 200 ms, which quite high. Smaller differences are not detected by the experiment.

The results for the help and distraction cases were expected. The highlighting helps, and it does not require attention. The last case to consider is the mixed case. The expected result was to see a combination of the help and distraction cases: decrease in the slope, and no change in the constant.

However, it was surprising to find that the mixed case does not have a significant slope decrease. The change in slope is -0.87 (s.e. 8.04, $t = -0.109$, $p = 0.914$) in the present case, and -11.29 (s.e. 8.84, $t = -1.28$, $p = 0.461$) in the absent case. This is much smaller than the help case results.

The only difference between the cases is the highlighting present on the current trial. Thus, the mixed case results indicate that information was successfully placed only into the user's iconic memory. When new information was presented, the iconic memory was overwritten. In effect, the user forgot where the highlighting appeared on the previous screen. Because the highlighting did not require attention to be useful, it means that information can be placed into iconic memory without requiring attention.

This experiment is promising, but it has two deficiencies. First, the test of distraction, how much time the user took to remember the position of the highlighting, is not sensitive enough. It is also not likely to be sensitive enough, because the amount of users needed to achieve high sensitivity on this task is quite large.

Second, the calm technology used in this study only gets the information to the iconic buffer. In order to conclude that an artifact does not require any attention, the information must be stored in the user's memory with no cost.

The objective of the rest of the chapter is to design an experiment that overcomes these deficiencies.

3.3 Methodology

To transfer information from an iconic buffer to memory, two steps are needed: encoding and storing. The encoding step typically requires attention. That is, if the user's attention is saturated by a task, then when information is presented, the user has two choices. He can delay encoding the information until attention is available. Essentially, this means the user transfers nothing into memory, because information does not stay in the sensory buffer for long. On the other hand, if the user chooses to encode information, by allocating attention to it, while performing the task, then performance on the task suffers.

Hence, the steps to design a dual task experiment that can detect attention required by encoding are the following.

1. Design a calm artifact to be tested. It must provide the user with non-trivial information, comprehension of which can be tested.
2. Design a primary task that saturates the attention necessary for encoding.
3. Show that the primary task actually saturates attention.
4. Present the artifact while the user is performing the primary task.
5. Measure performance changes on the primary task.
6. Measure how much information the user obtains from the artifact.

If the user acquires information from the artifact, without any degradation of performance on the primary task, then the calm artifact is unattended. However, if the user fails to acquire information, or if performance on the primary task decreases when it is presented, then the artifact requires attention for information acquisition.

3.4 Implementation

The rest of the chapter describes the details of the experiment. It begins with an overview of the information provided in the experiment and the artifact which is going to evaluated. Then, it describes the primary task along with cases that are designed to show that the task saturates attention. Finally, it describes the hardware and the procedure used during the experiment.

3.4.1 Information

The calm artifact provided information similar to that provided in previous calm technology research papers.

1. Temperature. The temperature was restricted to be one of three values: cold, temperate, and hot.
2. Time. The time was restricted to be one of four values: 6 o'clock, 9 o'clock, 12 o'clock, and 3 o'clock.
3. Time of Day. The time of day was restricted to be one of two values: day and night.

3.4.2 The Calm Artifact

The calm artifact is an abstract painting that forms the background for the primary task. It is referred to as the *background* in the remainder of the thesis (Figure 3.4).

Information Encoding

The information encoding was made intuitive to minimize the training time of users.

1. Temperature was represented as the dominant colour of the background. When the background was mostly blue, the temperature it represented was

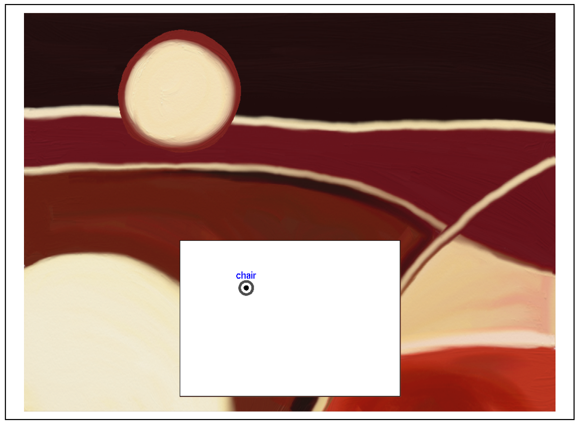

Figure 3.4: An example of the calm artifact (hot, 3 o'clock, day) as background for the primary task

cold. Mostly green meant the temperature was temperate, and mostly red meant the temperature was hot.

2. Time was represented by the position of a prominent circle on the horizon line, which was intended to represent movement of the sun from sunrise to sunset or the moon from sunset to sunrise. When the circle was on the left side of the screen, it represented 6 am/pm. Continuing from left side to the right side, it moved through 9, 12, and 3 in that order. Finally, when it was on the right side of the screen, it represented 6 pm/am.

Figure 3.5: An example of the background (cold, 6 o'clock, day) with the primary task

It is important to note that when the circle was at 6 o'clock, half of it appeared at the left side of the screen and half at the right side of the screen (Figure 3.5). That is why both positions represent 6 o'clock. This was done to ensure continuity of time representation. That is, going from 5:59 pm to 6:01 pm was not displayed as a sudden jump of the circle, but rather as a smooth transition. The circle on the right side of the screen moved outside the screen, thus becoming less visible, and circle of the left side moved closer to the center of the screen, thus becoming more visible.

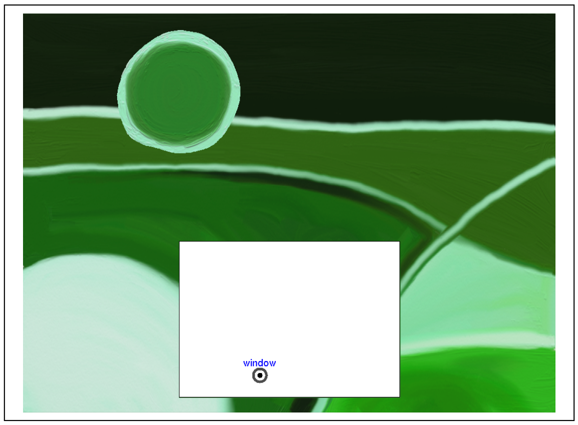

Figure 3.6: An example of the background (moderate, 9 o'clock, night) with the primary task

3. Time of day was signified by the circle on the horizon line, which was either a schematic sun or a schematic moon. The sun, shown as a bright circle with a dark rim, signified day. The moon, shown as a dark circle with a bright rim, signified night.

   When the circle is split, at 6 o'clock, half of it was the sun, and half the moon. For instance, at 6 pm, the half of the circle on the left side of the screen is the moon: night is falling. The half of the circle on the right side of the screen is the sun: day is ending. The user was instructed to pick the type of the circle on the right side of the screen as the correct answer. In Figure 3.5, the time of day at 6 o'clock is 'day'.

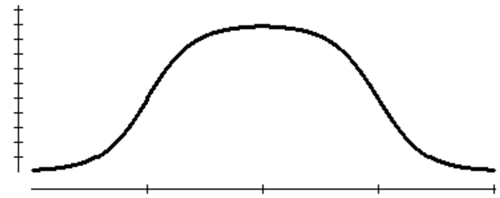

Figure 3.7: Fading function used

Presentation

The background was turned on and off using the smooth function plotted in Figure 3.7, a Gaussian. The function was chosen because all its derivatives are continuous.

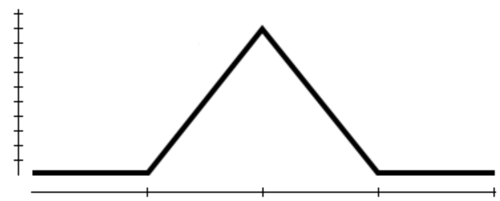

Figure 3.8: Linear fading function

It was chosen because I observed, while designing the experiment, that derivative discontinuities (Figure 3.8) caused involuntary allocation of attention to the background, in effect making it not calm. To minimize this effect, the fading function above was used, which ensured that the background appeared without forcing the user to allocate attention to it, and disappeared equally smoothly.

3.4.3 Gathering User Responses

To estimate how well the user absorbed information, a set of three multiple-choice questions was presented at the end of each trial. The answers were converted into percent correct values for each user and each question, which were, in turn, analyzed to determine how well the users absorbed information.

The questions ask about information presented during the trial. There were three questions: "temperature?", "time?", and "day/night?". The choices for each of the questions were the possible values of the information. For instance, for the question "temperature?", the possible answers were "cold," "moderate," and "hot".

The questions were asked by displaying a question such as "temperature?" and presenting a number of possible answers. For the temperature questions, the choices for the answer were "cold," "temperate," and "hot." The user had to select his answer by pressing 1, 2, or 3.

3.4.4 Primary Task

The primary task had three elements: two are tasks for the user to perform, and one is an element designed to further saturate attention.

The tasks that the user needs to perform during the trial are shadowing, reading aloud of simple words that appeared on the screen, and tracking a target with a mouse. An additional element that is present during the trial is the continuous background sound of random sentences in headphones, to further saturate attention.

Shadowing

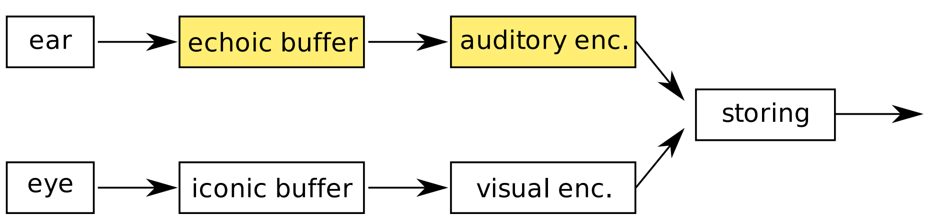

Figure 3.9: The shadowing task saturates a part of the memory process. It occupies the limited resources that would otherwise be used for auditory encoding.

Shadowing is commonly used to saturate the attention required for auditory encoding (Figure 3.9). When the user is occupied with shadowing, he cannot encode other auditory information, because the limited resources needed are not available.

Dolch's list of the most common English nouns was selected as the source of words to be shadowed. Two words were excluded, 'Santa Claus' and 'Christmas', because they were judged to be proper nouns.

During a trial, the words appeared at a rate of one every 1.5 seconds.

The user's responses were recorded, and later converted to *response times*. A response time is the amount of time that elapses from a particular word's appearance on the screen to the time when the user began reading the word, measured in milliseconds.

During the experiment, the recording of user responses was done using the OpenAL library. The conversion was done using the Audacity silence finder filter, followed by auditory inspection to ensure that no false positives were present. In other words, I found the beginning of each recorded word using the silence finder filter, and made sure that this was truly the beginning of the word being shadowed, not an unrelated sound, by listening to the recording.

Target Tracking

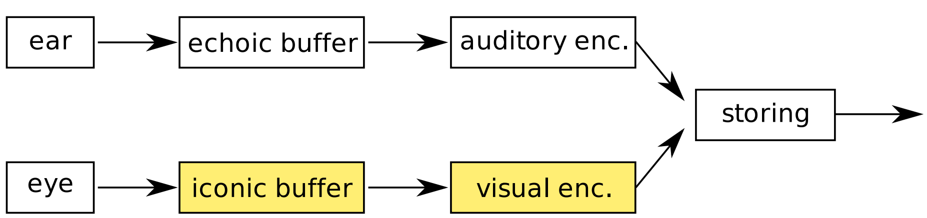

Figure 3.10: The tracking task saturates a part of the memory process. It occupies the limited resources that would otherwise be used for visual encoding.

Target tracking was implemented to saturate attention required for visual encoding (Figure 3.10). The speed of the target was adjusted for each user at the beginning of the experiment, so that the user could only follow it successfully 80% of the time, defined as the amount of time the mouse cursor was within 20 pixels of the target's center. This was done so it would be easy to determine if resources available to the tracking task were diminished. If the user was tracking the target perfectly, then visual encoding resources might not be 100% saturated. Hence, the user following the target successfully 80% of the time is a good base case.

The *success rate* for target tracking was measured in order to detect any attention claims that caused the user to perform poorly. It was calculated by counting the number of frames with the mouse position within 20 pixels of the target's center. This number was converted into a percentage and logged.

Playback of Random Sentences

The playback of random sentences was added to present the user from using the echoic buffer for short-term memory. The random sentences were selected from National Geographic, minimizing any particular topic. That is, the sentences were not related to the experiment, computers, each other, or any other particular topic.

I recorded the sentences in a neutral voice, and played them in a random order one after the other during the experiment. This yielded a continuous track of random sentences.

3.4.5 Testing Attention Saturation

To show that the primary task saturates attention, the background was replaced by other information sources, which required amounts of encoding comparable to the background. That is, the background no longer appeared during the trial and information was provided using other artifactd, which were known to be attended. Since they required attention, then it is easy to determine whether the primary task saturated the user's attention. If the user could not acquire information from the artifacts without disturbing the primary task, then the primary task saturated attention.

The artifacts which were designed to require attention are discussed below.

Aural information

The artifact that was used during the experiment to show that auditory encoding was saturated was an audio recording. That is, the information was presented in words played into the user's headphones. The recordings, which contained the same information as the background, were placed between the random sentences to create a continuous audio track.

A set of audio recordings was made, including sentences like "The time is six o'clock," "The temperature is hot," and "It is daytime." These were placed at predefined intervals on the track of random sentences.

The intervals were picked to be spread throughout the trial. Temperature information was played at one fourth of the total trial time from the beginning. Time information was played half way through the trial. Time of day information was played at one fourth of the total trial time from the end.

When each sentences was played varied from trial to trial. Since the information was integrated seamlessly among the random sentences, it was only possible to play the information at an *earliest possible* time, at the end of the random sentence.

This timing ensured two things. Firstly, the information was seamlessly integrated into the audio track. This was important, because the playback of information did not force the user to allocate attention to it. Secondly, the beginning times of information playback were jittered. This minimized the possibility of a user performing well or poorly because of a particular instant the information playback began.

Control Experiment

To show that the primary task saturates attention used for visual encoding, an informal control experiment was conducted with three users.

The same information was provided using the background that, instead of fading in and out, appeared suddenly half way through the trial. This ensured that the user allocated attention to the background.

3.5 Trial Details

Each trial lasted 90 seconds. There were a total of 24 trials per experiment, 12 of the aural type and 12 of the background type.

Every permutation of information was used once during each experiment. This means that each trial had a different set of information. This ensured that the user did not access their long term memory in order to answer the multiple-choice questions that followed. Also, it helped minimize the effect of information combinations that were uncommonly easy or difficult. For instance, hot temperature might be easily associated with 12 noon, in which case repeating the combination would bias the correctness results. By ensuring no case was repeated, this bias was eliminated.

Together with training and breaks, each user spent approximately 50 minutes on the experiment.

3.5.1 Hardware

The experiment was conducted on a typical desktop machine:

- Pentium 4 CPU, 2.4 gHz,
- ATI Radeon X1900 graphics card,
- 19" LCD monitor, running at its native 1024x768 resolution, and
- Standard keyboard, mouse, and headset

3.5.2 Procedure

The user was seated approximately 70 cm from the monitor, and was able to adjust his seating until he was within comfortable reach of both the mouse and the keyboard. The headset was also adjusted to be comfortable.

The primary task, the aural case, and background cases were fully explained to the user using a script as a guide. The information encoding in the background was explained in detail. Furthermore, the user was told that at the end of each trial there would be a set of multiple-choice questions enquiring about the trial just ended.

The user was allowed to perform several sample trials to become accustomed to the primary task. Afterwards, the user had a chance to ask questions about any aspect of the experiment that was not clear to him. Finally, the user was left alone in a quiet dimly lit experiment room to perform the experiment.

After each trial, there was a short break. The user was instructed that this was a short break, but should not last more than a few of seconds. There were also two long breaks spaced out evenly in the session. The long breaks allowed the user to take a couple of minutes to relax, so that the performance on the experiment would not degrade towards the end due to fatigue.

Users

The users were recruited using posters, mailing lists, and word of mouth. A total of seven users performed the experiment. All users were in their early 20s. Two users were female.

The users received a bottle of water to drink during the experiment, and $10 upon completion.

After all users had performed the experiment, the results were analyzed and interpreted. The results of the experiment are the topic of the next chapter.

Chapter 4

Results

This chapter describes the results of the experiment and provides low level interpretation of the results. The next chapter contains further interpretation.

4.1 Layout of the Argument

The experiment investigates whether attention is required for information encoding. In particular, it attempts to show that attention is not required to acquire information from the background.

The primary task consists of two parts: shadowing and target tracking. Shadowing is limited by one of two things: by auditory encoding or by speech comprehension and generation. That is, the user cannot perform better either because resources required for auditory encoding are occupied or because the user cannot produce speech fast enough. Similarly, tracking is limited either by resources required for visual encoding or by motor response.

It is essential to establish that the primary task saturates attention, because it needs to detect attention demands that are made elsewhere. The results of the background evaluation show that information provided by the background does not disturb the user's performance. If the primary task saturates attention, it means

user number	temperature	time	time of day
1	100%	91.7%	100%
2	83.3%	83.3%	91.7%
3	100%	91.7%	75%
4	83.3%	100%	83.3%
5	91.7%	91.7%	91.7%
6	58.3%	75%	83.3%
7	100%	91.7%	100%
average	88.1%	89.2%	89.2%

Table 4.1: Success rate for aural case multiple-choice responses.

that the background does not require attention to provide information. Hence, the first part of the chapter presents the results that establish that the primary task saturates attention needed for both visual and auditory encoding.

The second part of the chapter discusses the results of the background case, which show that the background provides information without requiring attention. Hence, it shows that the background is unattended.

4.2 Saturation of Attention

In order to show that the primary task saturates the attention needed for auditory encoding, results of both the aural case and the control experiment need to be examined.

4.2.1 Aural Case Results

Multiple-choice Correctness

From the average success rates, it is evident that users answered the questions well, but far from perfectly (Table 4.1). On average, only 88-89% of the questions were answered successfully. Since the trials were 90 seconds in length, this is a low success rate.

This is one indication that the user's attention was saturated. Most of the time, the user was able to encode information, but the rest of the time, he was unable to encode the information while maintaining satisfactory performance on the primary task. Hence, information was not stored in memory and the multiple-choice success rate decreased.

The success rate is similar across all users, suggesting that any set of users would produce similar results.

Shadowing Response Times

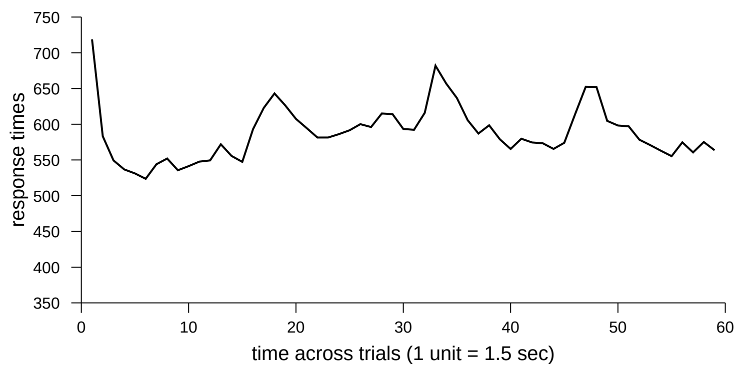

Figure 4.1: Average response times for the aural case.

There are four visible increases present in the response times graph (Figure 4.1).

The initial increase (t=0) is a consequence of the experiment design. First word is unexpected, and it takes a word or two for the user to get into the rhythm of speaking.

The other three increases occur at exactly the time that information was provided aurally. These are the important increases to investigate, since they show that encoding the information requires attention, taken from shadowing.

The times at which information was given were jittered from trial to trial. Aligning the start of information across the trials before averaging shows a consistent

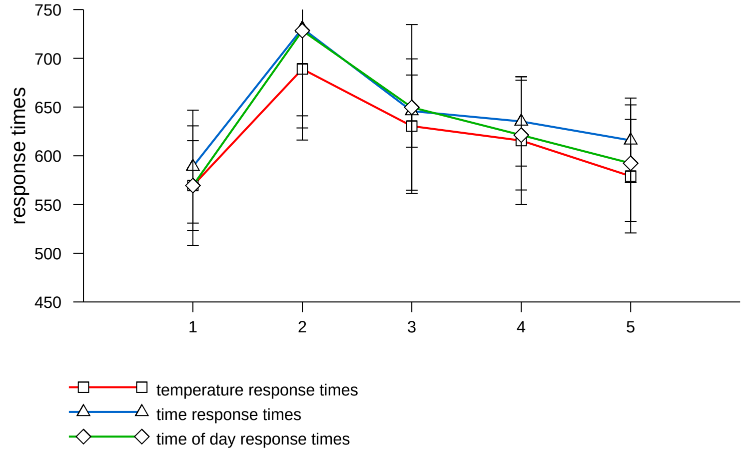

Figure 4.2: Aligned average response times for the aural case.

pattern of increases (Figure 4.2), indicating that the resources taken from shadowing follow a constant pattern.

The increases indicate that shadowing is limited by the attention required for auditory encoding. Since presented information did not need to be shadowed, yet increases in the response times appeared, the shadowing task is not limited by language generation. The only other explanation is that shadowing is limited by attention. Hence, shadowing saturates the attention required for auditory encoding.

Tracking Task Performance

Results of the tracking task reveal that tracking was not disturbed by the aural case (Figure 4.3). The graph indicates that users were consistent, somewhat below perfect, in the tracking task throughout the trial.

Since the attention needed for auditory encoding was already saturated by the shadowing task, this result confirms that attention allocated during auditory encoding is separate from attention used for visual encoding, in accord with the results of Baddeley and Hitch [2].

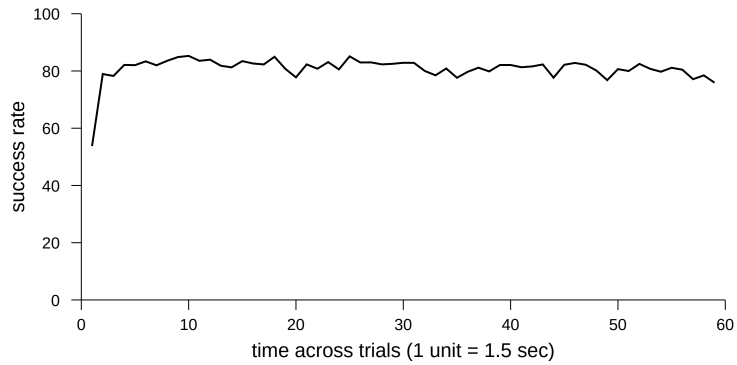

Figure 4.3: Average success rates for the aural case, showing the start of aural information at time 1 and the following four points.

4.2.2 Control experiment

Since it is established that the primary task saturated the attention available for auditory encoding, the only results discussed for the control experiment are the target tracking results.

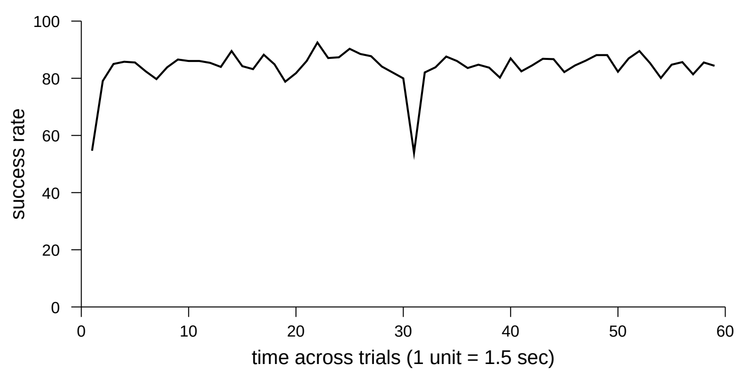

Figure 4.4: Average success rates for the control experiment.

The tracking task graph (Figure 4.4) reveals that there are two decreases on performance that occurred during the trials. The initial decrease can be accounted for by experiment design. It occurs because the target started moving abruptly, and it took users small amount of time to get used to its movement.

The second decrease occurs in the middle of the trial, exactly when the background suddenly appeared. Hence, the onset of the background diminished the user's tracking performance.

It is well-known that sudden changes in visual stimuli automatically claim attention, and because this disruption was registered as a decrease in tracking performance, target tracking is limited by attention. That is, visual encoding occupied by the tracking task.

4.2.3 Interpretation

The aural case results show that the performance on the shadowing task is limited by the attention required for auditory encoding. That is, when presented with aural information, the user's performance on the task suffered. This is clear from the increases in the aural case's response time graphs.

Furthermore, the user was not able to acquire information in all trials. The low success rate indicates that part of the time, the user chose to perform the shadowing task, instead of encoding and storing aural information. This confirms that the shadowing task saturates attention needed for auditory encoding.

Visual encoding is also saturated. This is evident in the results of the control experiment. Target tracking results were disturbed when the background appeared suddenly. This indicates that the tracking task is limited by the attention required for visual encoding. In turn, this means that this attention was saturated.

These results show that the primary task saturates attention. It can be used in a dual task experiment to measure whether or not attention is required by an artifact. In fact, I have done this to see whether the calm artifact I designed is unattended.

user number	temperature	time	time of day
1	100%	91.7%	91.7%
2	100%	100%	100%
3	100%	100%	100%
4	100%	100%	100%
5	100%	91.7%	100%
6	100%	100%	100%
7	100%	100%	100%
average	100%	97.6%	98.8%

Table 4.2: Success rate for background case responses.

4.3 The Background Evaluation

4.3.1 Correctness

The background shows a significant improvement in the multiple-choice success rates (Table 4.2) compared to the aural case. The averages are close to 100%, indicating that the background provides information well.

4.3.2 Response times

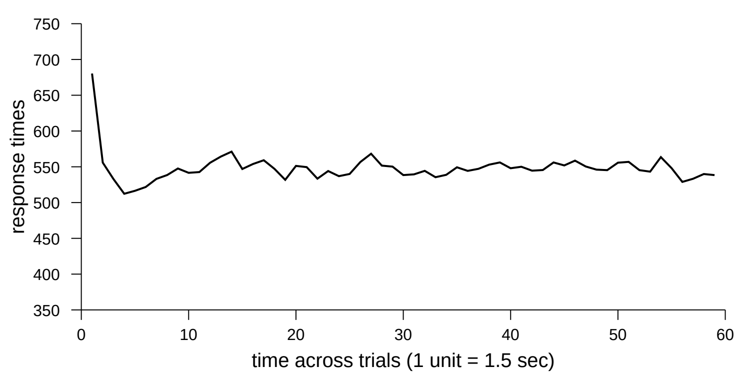

Figure 4.5: Average response times for the background case.

It is important to check response times for the background case, because attention might be required to encode information linguistically for long-term storage.

However, the response times graph indicates that, aside from the initial increase, there are no significant increases or decreases (Figure 4.5).

This means that attention required for auditory encoding is not used. Of course, the background could require visual attention, not auditory. However, the fact that shadowing response times are not affected is necessary to conclude that the background does not require any attention.

4.3.3 Target tracking performance

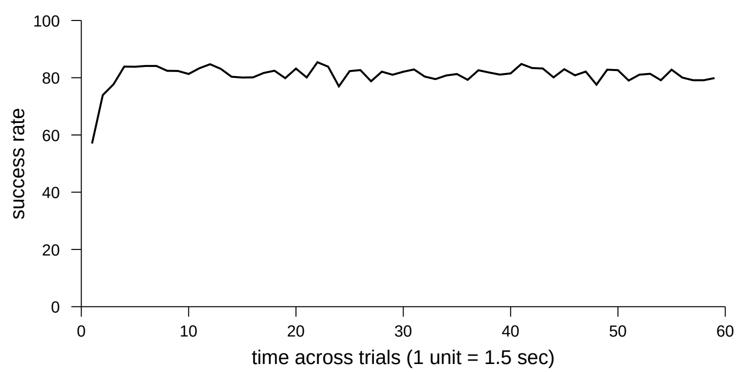

Figure 4.6: Average success rates for the background case.

Target tracking is also unaffected by the background (Figure 4.6). The result is important for the background case, because if it required attention for visual encoding, the target tracking would be affected.

4.3.4 Interpretation

The background results show that there are no disruptions in either the tracking or the shadowing tasks. Furthermore, information acquisition is essentially perfect.

Since the primary task saturates attention, the background results indicate that it does not require attention auditory or visual encoding. This means that the background is unattended.

The multiple-choice success rate is near perfect. This means that the user successfully acquired information from the background. Thus, the background provide information without requiring attention.

This shows that unattended technology is possible: I have provided an example of it that possesses its identifying features.

Chapter 5

Conclusion

Initially, this chapter interprets the results of the experiment. It reveals that the background is unattended, providing information without requiring attention.

Afterwards, the chapter considers some possible future research. As well, concluding thoughts are presented.

5.1 Interpretation of the Results

This thesis makes three contributions to calm technology research. First, it defines unattended technology as the limiting case of calm technology. That is, unattended technology is calm technology that does not require any attention when providing information. Second, this thesis develops a method for the evaluation of unattended technology. Using a dual task experiment with an attention saturating primary task, it is possible to determine whether a calm artifact is unattended. Last, this thesis shows that unattended technology is possible by showing that a particular example of it passes evaluation.

The definition of unattended technology makes it possible to talk about this subset of calm technology without any ambiguity. Calm technology's definition includes the three criteria listed in chapter 2. The last criterion of calm technology states that calm technology has to provide information while requiring little attention, which is fairly vague: how much is 'little'? However, unattended technology is precise: it requires no attention while providing information.

Evaluation of unattended technology is also important. With an evaluation procedure we can look at new and existing calm artifacts and determine whether or not they are unattended. That is, it is now possible, at least in principle, to determine whether Weiser's LiveWire, for instance, requires more attention than Cadiz et al.'s Sideshow. In effect, researchers are now able to make concrete claims about their calm artifacts. For instance, it is possible to say that an artifact requires less attention than LiveWire, but does require some attention, so that it is less attended, but not unattended.

Finally, showing that unattended technology is possible is a significant contribution. It shows that a user can acquire information while expending no measurable attention. This provides an existence proof, which is a first step to scientifically investigating unattended technology. To get a clear understanding of the limits of unattended technology, more research is needed.

5.2 Limitations of the Experiment

The experiment conducted in this thesis produces results showing that the background is unattended. The results are based on the conclusion that the primary task fully saturates the attention needed for auditory and visual encoding. However, there might be a different interpretation to the results. In particular, target tracking might not fully saturate attention needed for visual encoding.

Consider the possibility that target tracking is limited by motor response generation and it requires most, but not all, of the attention available for visual encoding. Making the background appear suddenly during the control experiment might require more of the visual attention than is available. That is, the user's attention required for visual encoding is overloaded, because the abrupt background onset claims much attention. Hence, the result is a decrease in performance when the user is exposed to the background. However, presenting the background using a fading function might also require attention, but it does not overload it.

This means that the results show that the background requires much less attention than other, attended, artifacts. However, this does not necessarily mean that the background requires no attention. In order to provide more evidence of this, more research is needed.

5.3 Future Work

This thesis provides an existence proof for a limiting case of calm technology, unattended technology, and is a first step in exploring calm and unattended technology rigorously. However, more work will be needed to reach a full understanding of calm technology. Three suggested research directions suggestions are described below.

5.3.1 Improving the Primary Task

To show that unattended technology does not require attention, the null hypothesis must be proved. I have shown that while the primary task I designed can detect some level of attention claims, it remains possible that small attention claims are made, but were not detected.

To provide stronger evidence for the existence of unattended technology, attention needs to be quantified. Furthermore, the primary task needs to be improved so that more concrete claims can be made about the attention demand of unattended artifacts. By measuring attention in precise units, one will be able to claim that the primary task would be able to detect a certain quantity of attention demand. In turn, this would allow a conclusion to be drawn that an unattended artifact requires less attention that this quantity. By improving the primary task, one will be able to detect smaller quantities of attention claimed.

This can be done by either improving my primary task, or introducing a new and better primary task. The primary task can be improved by fine tuning the parameters, such as the speed at which words to be shadowed appear. New tasks that can be introduced include known attention saturating tasks, such as the search task I used for the pilot study (Ls and Ts).

5.3.2 Measuring the Amount of Information

I have shown that unattended technology can provide a small amount of information, without trying to measure how much information can be reliably provided.

To use unattended artifacts effectively, it is necessary to know how much information they can reliably provide. This can be done by gradually increasing the amount of information until either an attention requirement appears, or the information is incompletely acquired.

Such results would show how much information can be provided by a single unattended artifact.

5.3.3 Measuring Whether Information is Precise

Furthermore, the information provided by the unattended artifact is of low granularity. For instance, the time is restricted to be one of only four values. For most purposes this is not precise enough to be useful. Ideally, the user should be able to tell time with at least a half-hour precision. For example, when working in an office, the user should be able to determine whether it is lunch hour, not whether the lunch hour is at most three hours from now.

To find out how well unattended artifacts can provide precise information, the resolution of information could be gradually increasing, while changes in primary task performance are measured. This would reveal how precisely unattended artifacts can provide information.

5.3.4 Other Properties

Other properties of unattended technology need to be investigated. In the experiment I designed, the unattended artifact lies all around the screen space occupied by the primary task. How would the attention demands of information acquisition change if the unattended artifact were moved away from the primary task? In practice, unattended artifacts are unlikely to be near the center of the user's visual field, making this an important question to answer. Hence, investigating how effectiveness varies with position in the visual field is important.

This thesis is a first step in exploring unattended technology, and properties of calm technology as they relate to attention demand. More research is needed before we can construct an attention profile of an unattended artifact.

5.4 Discussion

The main contribution of this thesis is to show that information can be provided without requiring the user's attention, that unattended technology is possible.

I show that unattended artifacts exist. This means that they can be introduced without requiring attention from the user: it may be possible to have a lot of information available without spending time and attention acquiring it.

This has deep implications. Now, non-trivial information costs attention. It is a simple fact, but it means that to learn any piece of information, one has to spend attention. But with unattended technology, one can absorb information without requiring attention. This means that one can spend attention on other matters, while getting constant updates on weather conditions, for example.

If it is discovered that unattended technology can provide more information, and precise information, then unattended technology not only *can* be placed into homes and offices, it *should* be placed there. It will change from being a novelty, a toy, to an important part of our life. By providing information without requiring attention, unattended technology essentially gives the user the capability of doing more tasks, than he would be able to without unattended technology. For instance, reading a text while frequently checking the time is both tiring and distracting. However, if time is provided using unattended technology, then the user can concentrate on reading and without losing any time information.

The first step is investigating the concrete properties of calm technology. This will result in an understanding of what calm technology can and cannot accomplish.

In all probability, unattended technology cannot provide as precise information as attended technology can. However, this does not mean that unattended technology is not useful. When working at its full potential, unattended technology should provide information to the user, while being undetectable to the user. For instance, a key press on the keyboard provides tactile and force feedback to the user. This information is small, but important. It allows the user to type a password, for example, and feel confident that it is correct. However, the user rarely notices or pays attention to the feedback that keys provide. This exemplifies how unattended technology can be useful.

Consider another example. Suppose the mechanical repair of a car requires a set of tools to be used in a particular order. Instead of memorizing the tool order, or referring frequently to a manual, the mechanic could install an unattended artifact to help him with the task. The unattended artifact would highlight the next tool needed on a work bench, letting the mechanic know to reach for it. Ideally, after some time, the mechanic would neither remember which tool is needed nor refer to the manual nor remember that an unattended artifact is helping him. He would simply reach for the right tool all the time.

Such a goal seems too good to be true. However, the unattended bandwidth of sensory systems is vast, especially compared to the bandwidth of consciousness. Unconscious processing allows us to perform many seemingly simple tasks, such as sitting in a chair without falling, walking in a straight line without tripping, or reacting to unsafe drivers on the road, are all possible with minimal attended processing. Unattended technology tries to provide non-trivial information through the same channel. If it succeeds in practice, then some tasks that currently require attention could be done as easily as reacting to a stop sign by pressing the brakes.

Unattended technology can alleviate the stress of attention-demanding tasks. It can provide information without requiring attention, and without distracting the user. In essence, while the user is under the impression that he is just 'that good' at performing a task, unattended technology is providing the information necessary for the user to complete the task.

References

[1] R.C. Atkinson and R.M. Shiffrin. Human memory: A proposed system and its control processes. *The psychology of learning and motivation*, 8, 1968. 22

[2] A.D. Baddeley and G.J. Hitch. Working memory. *The psychology of learning and motivation: advances in research and theory*, 8:47–89, 1974. 22, 42

[3] J.J. Cadiz, G. Venolia, G. Jancke, and A. Gupta. Designing and deploying an information awareness interface. In *Proceedings of the 2002 ACM conference of Computer supported cooperative work.*, pages 314–323, 2002. 8, 9, 14, 16

[4] S. Consolvo and J. Towle. Evaluating an ambient display for the home. In *CHI '05 extended abstracts on Human factors in computing systems.*, pages 1304–1307, 2005. 8, 14, 16, 17

[5] E. Costanza, S.A. Inverso, E. Pavlov, R. Allen, and P. Maes. eye-q: Eyeglass peripheral display for subtle intimate notifications. In *Proceedings of the 8th conference on Human-computer interaction with mobile devices and services.*, pages 211–218, 2006. 10, 18

[6] J.J. Gibson. *The Senses Considered As Perceptual Systems.* Houghton Mifflin, 1966.

[7] J.M. Heiner, S.E. Hudson, and K. Tanaka. The information percolator: Ambient information display in a decorative object. In *Symposium on User Interface Software and Technology.*, 1999. 8

[8] L.E. Holmquist. Evaluating the comprehension of ambient displays. In *CHI '04 extended abstracts on Human factors in computing systems.*, pages 1545–1545, 2004. 18

[9] L.E. Holmquist and T. Skog. Informative art: information visualization in everyday environments. In *Proceedings of the 1st international conference on Computer graphics and interactive techniques in Australasia and South East Asia.*, pages 229–235, 2003. 8, 9

[10] R. Iqbal, J. Sturm, O. Kulyk, J. Wang, and J. Terken. User-centred design and evaluation of ubiquitous services. In *Proceedings of the 23rd annual international conference on Design of communication: documenting & designing for pervasive information.*, pages 138–145, 2005. 8, 10, 14

[11] N. Jafarinaimi, J. Forlizzi, A. Hurst, and J. Zimmerman. Breakaway: an ambient display designed to change human behavior. In *CHI '05 extended abstracts on Human factors in computing systems.*, pages 1945–1948, 2005. 8, 10, 16, 17

[12] H. Lieberman. Autonomous interface agents. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 67–74, 1997.

[13] B. MacIntyre, E.D. Mynatt, S. Voida, K.M. Hansen, J. Tullio, and G.M. Corso. Support for multitasking and background awareness using interactive peripheral displays. In *Proceedings of the 14th annual ACM symposium on User interface software and technology.*, pages 41–50, 2001. 10

[14] P. Maes. Agents that reduce work and information overload. *Communications of the ACM*, 37(7):30–40, 1994.

[15] P.P. Maglio and C.S. Campbell. Tradeoffs in displaying peripheral information. In *Proceedings of the SIGCHI conference on Human factors in computing systems.*, pages 241–248, 2000. 18, 19

[16] J. Mankoff, A.K. Dey, G. Hsieh, J. Kientz, S. Lederer, and M. Ames. Heuristic evaluation of ambient displays. In *Conference on Human factors in computing systems.*, pages 169–176, 2003. 14

[17] D.S. McCrickard, C.M. Chewar, J.P. Somervell, and A. Ndiwalana. A model for notification systems evaluation assessing user goals for multitasking activity. In *ACM Transactions on Computer-Human Interaction (TOCHI) archive.*, pages 312–338, 2003. 10

[18] Herbert Simon. *Computers, Communications and the Public Interest*. The John Hopkins Press, 1971.

[19] T. Skog. Activity wallpaper: ambient visualization of activity information. In *Proceedings of the 5th conference on Designing interactive systems: processes, practices, methods, and techniques.*, pages 325–328, 2004. 8, 9, 10

[20] H. Stromberg, V. Pirttila, and V. Ikonen. Interactive scenarios: building ubiquitous computing concepts in the spirit of participatory design. *Personal and Ubiquitous Computing*, 8(3):200–207, 2004. 14

[21] M. Tomitsch, K. Kappel, A. Lehner, and T. Grechenig. Towards a taxonomy for ambient information systems. In *Ambient Information Systems*, 2007.

[22] W. van der Hoog, I. Keller, and P. Strappers. Gustbowl: technology supporting effective communication through routine ritual interactions. In *CHI '04 extended abstracts on Human factors in computing systems.*, page 775, 2004. 8, 16, 17

[23] D. Vogel and R. Balakrishnan. Interactive public ambient displays: transitioning from implicit to explicit, public to personal, interaction with multiple users. In *Proceedings of the 17th annual ACM symposium on User interface software and technology.*, pages 137–146, 2004. 8, 10, 11

[24] M. Weiser and J. Brown. Designing calm technology. PowerGrid Journal, v1.01 (see: http://www.powergrid.com/1.01/calmtech.html), 1996. 6, 8, 10