Context aware interruptions: existing research and required research

Ari Kolbeinsson, Peter Thorvald and Jessica Lindblom

School of Engineering Science
and School of Informatics
University of Skövde
SE 541-28, Sweden

ABSTRACT

Companies in various fields are developing information systems that are supposed to show information in a context sensitive manner. This involves modifying user interfaces based on the current activity and location of the user. The research that exists for context aware, task aware, and location aware systems come mostly from research in the office environment and using stationary work computers using standard mouse and keyboard-operated systems. Other environments, such as manufacturing environments, have not received the same research attention, and therefore research is required to see whether existing theories and frameworks apply to the manufacturing domain. Adding to this, context aware systems are now being created for new classes of devices such as mobile and wearable devices, to be used in multiple domains. All this requires an investigation and validation of older research, and shows how the research of the basic human factors surrounding new devices and domains has fallen behind the development of the devices themselves. This paper examines shortly how recent changes advances in technology affect what is required from the field of interruption research, as well as what is needed to support other domains than the office environment.

Keywords: Interruptions, interruption management, context aware, user interfaces, manufacturing

INTRODUCTION

The use of mobile devices has increased greatly in recent years, with these devices having become near ubiquitous in parts of the world. These handheld devices make possible instant communications at all times, as well as automated notifications of all sorts. This happens during people's time off, during working hours, during their commute, and while they sleep. Some of these devices are privately owned and used for personal communications, while others are owned by the workplace and are in some cases specially designed for supporting information services within the company. Yet these devices, and commonly available consumer devices in particular, mostly fail to take into account the cost of interruptions. Interruptions cause a loss in effectiveness on the task being performed when the interruption strikes, the primary task, and this loss in effectiveness may range between annoying to dangerous, dependent on the criticality of the primary task (McFarlane and Latorella, 2002). The interruption task may also have its own requirements. Some interruptions can be critical, others may require response of some sort in a timely manner, but most interruptions can wait for a short time. Current mobile devices are now being expanded through the introduction of wearable devices, which are finally coming to the mass market, but these suffer from the same limitations as mobile devices in that interruptions are not effectively managed.

All this raises a question: why are all these new devices using decades old notification/interruption methods?

The answer to that lies in the current state of research, especially research on notifications/interruptions.

Interruptions have been formally researched since the late 1920's when it was discovered that people who were interrupted from their primary task could better recall the details of the interrupted primary task than the details of task that was not interrupted (Spiekermann and Romanow, 2008). Research on interruptions then increased with the rise in research on human factors after some highly publicised disasters in fields such as aviation and nuclear power stations.

When examining existing research on interruptions it becomes clear that there is a focus on information workers in an office environment, which has been both the area where computer systems have been most likely to be used as well as an area that is relatively easy for researcher to gain access to or simulate. Offices have thus received much attention, and the stationary desktop machines that have been the staple of offices have likewise received the majority of attention. Bjelica, Mrazovac, Papp, and Teslic (2013) state that research into interruptions in the field of human-computer interaction (HCI) and context-aware computing only started with McFarlane (2002), which seems surprisingly recent. Earlier examples can be found in specialist domains such as the Fokker 50 aircraft cockpit where systems changed interruptions based on the current requirements as early as 1987 (Fokker, 2012), but these were not referred to as *context aware* systems at the time, and are highly specialised.

Keeping the focus only on the office environment is not a strategy that supports current or future uses of technology, as the landscape of information services and information dense environments has changed in the last decade. Miniaturisation of computing devices affects more than just the size of the device; miniaturisation fundamentally changes the possible context of use from being stationary at a desk or standing at a workstation to being anywhere and in any position. This changes how users can interact with the device and places new cognitive demands on the user; demands that require research that is specifically aimed at examining the effects of interruptions by mobile devices on users' cognitive facilities. As mobile devices become more prevalent it becomes necessary to examine the effect of those, and this should preferably be done in multiple settings as mobile devices are permeating almost all aspects of work and play.

This change due to mobile devices also reaches to many workplaces. Work has increasingly become non-linear and non-exclusive, meaning that people have to hop between tasks and juggle multiple tasks at a time, even juggling information between multiple computing devices. Office environments have been designed at least in part around this, and office computers, both desktop computers and notebook computers, incorporate task-switching mechanisms (Speier, Vessey, and Valaich, 2003). Handheld devices are only partly supported, even in the office environment, and mostly do not cooperate with the desktop systems. The quick change in information services due to mobile devices is clearly portrayed by Speier et al. (2003) who mention mobile devices only being a source of interruptions due to ill-timed or unwanted phone calls.

Manufacturing environments have not had the same level of information flow to workers, which stemmed from production environments focusing on mass production of series of identical units. With current requirements of product diversity, customization, and quality control it has become more important than before to effectively communicate changes in products to line workers. This requires either paper manuals, which must be manually retrieved from some sort of storage area and is inefficient in a quickly changing product line, or a stationary computerised solution. Foremen must also respond effectively to quickly changing conditions and all manner of problems, and so can be supported by an effective mobile solution to replace less technical solutions such as shouting or searching for a foreman or using a klaxon or a blinking light at the workstation.

This suggests some technical solutions that are in many ways simple to implement, but there are many other factors that must be solved to make an effective and agreeable solutions. While being instantly notified of problems is an effective for the message, it may not be effective for the task that the receiver of the message is working on at the time. Current systems offer audio, video, and vibration as notification signals while stationary devices can also use connected peripherals, and future systems such as head mounted heads up displays or augmented reality displays may incorporate all of these in some manner. One of the problems faced is then one of how interruptions by information systems affect workers in non-office environments, such as a manufacturing environment.

What Is Meant By "Interruption"?

An interruption is any event that breaks a person's attention to their current activity (primary task) and either requests or forces the user to focus on a new task (interruption task). Corraggio (1990, p. 19) provides a definition: "an interruption is an externally generated, randomly occurring, discrete event that breaks continuity of cognitive focus on a primary task". Interruptions can cause increases in errors on the primary task and induce errors on the interruption task, but McFarlane and Latorella (2002) argue that this can be minimized through taking into consideration the requirements of the person, the primary task, and the interruption task.

What Is Meant By "Notification"?

The terms *notification* and *notification system*, as used here, come from computer system makers who have called their interruption systems *notification systems* (Apple, 2012; Microsoft, 2012). Although notifications are widely discussed in research, a clear definition of what a notification or a notification system is was not found in the literature, which is why a definition is proposed here. A notification is the product of an interruption mechanism (the notification system) that interrupts a person's current (primary) task and provides information that another task requires attention. The notifications themselves can be information rich or information sparse. A simple bell sound from a mobile device may notify the user that a message has been received, but gives no information about the content of the message or whether this message requires immediate attention. This paper refers only to notifications from information systems, which use systems that are specifically designed to interrupt the user. Notification system design has been identified as a major component in mitigating the costs of interruptions when using information systems (McFarlane and Latorella, 2002)

What Is Meant By "Context Awareness"?

The words *context aware user interface* can cause some confusion due to their being used for multiple purposes. This paper discusses context in its widest meaning. A *context* is: "the circumstances that form the setting for an event, statement, or idea, and in terms of which it can be fully understood" (New Oxford American Dictionary, 2005). An early definition of *context aware computing* comes from Schilit, Adams, and Want (1994) who explain that context aware systems adapt according to multiple factors in the system's environment, such as the location, people present, other devices, and changes over time. Schilit et al. (1994) further explain that *context* encompasses such things as the sound and light levels, communication bandwidth, and even the social context, i.e. *who* may be present. This shows *context* as being a very wide concept that can include almost whatever the designer wants, but it also suggests that any system billing itself as being *context aware* should take into account all factors that affect the task at hand, or a clearly defined subset of those factors. This means that just making a UI that changes does not mean that it is a context aware system. The UI has to change and adapt to the users' needs and abilities, and cognitive abilities have to be taken into account from the very beginning. This is where many systems fail. Some of the systems discussed in this paper, for instance, take into account technical needs and in some cases also what the designers think that the user might *want* to see, but fail to take into account what the user actually *needs* and what the user's cognitive facilities can actually *handle*. The lack of research that supports the design of context sensitive systems outside of well defined situations using desktop computers must be considered a factor in this.

THE CURRENT SITUATION IN RESEARCH

McFarlane and Latorella (2002) mention that it may seem absurd that people may actually want interruptions to their work, but that normal work requires access to updated information and recreational time can be well supported by access to social information. Research on interruptions is nothing new, with some of the oldest known research harking back to the late 1920's when it was found that a person who was working on one task (the primary task) and got interrupted by another task wound up remembering the details of the interrupted primary task better (Zeigarnik, 1927). This is referred to as the Zeigarnik effect (Spiekerman and Romanow, 2008). Research on the effect of interruptions then took off in the later half of the twentieth century (Spiekerman and Romanow, 2008), and with research increasing much with the rise of research into human factors and ergonomics (Hawkins, 1987). Yet McFarlane and Latorella (2002) also point out that information systems at the time of their research had in general not been designed with interruptions taken into account and that interruptions often have dangerous effects in workplaces such as aircraft cockpits where pilots were much more likely to make mistakes on interrupted tasks. In fact, research on interruptions was found mostly in psychological research until the mid 1970s' when interruption research started appearing in ergonomics. The 1980s' saw interruption research starting to appear in information

https://openaccess.cms-conferences.org/#/publications/book/978-1-4951-2101-2

systems research, , but the field grew immensely in the late 1990s' (Speikermann and Romanow, 2008). It may seem prudent to ask why information systems design did not take the effect of interruptions into account, but by now the answer has become apparent. Zijlstra, Roe, Leonora, and Krediet (1999) state that little research had been done by 1999 on the effects of interruptions on workers' performance, which is in line with Bjelica et al. (2013) who discussed the lack of research on interruptions in HCI and context aware computing up until 2002. Examining existing literature shows this to contain a grain of truth. Even though the earliest research stems from 1927, and research has expanded greatly after computers and information systems became a part of everyday life in the 1990s (Spiekermann and Romanow, 2008), there is still a lack of basic research on interruptions in most domains, and what research exists consists of standard psychological testing methods using paper cards, or uses stationary desktop computers. Both these approaches fail to take into account the dynamic nature of today's information landscape, where activities and information flow can coexist simultaneously in all forms of work or recreation. This can be seen in work such as Altman and Trafton (2007) and Jones, Gould, and Cox (2012) that use stationary computers, as well as in Warnock, McGee-Lennon, and Brewster (2011) and Morgan and Patrick (2013) that use puzzle based psychological research methods that do not reflect real world situations and may therefore lack ecological validity.

There are, of course, exceptions. Some organisations such as the US Navy are reported as having become aware quite early that the complexity of their shipborne systems was increasing in such a way that future systems would not be fully mission capable in the future without a major UI redesign that would take interruptions into account in notification systems (McFarlane and Latorella, 2002). Although this example uses workplaces with large consequences for mistakes, the same is likely to apply to other fields. Even though the consequences of mistakes may not be fatal, they may still be harmful or even disastrous to the individual or company that suffer these problems. Yet other mistakes may lower the operational efficiency of a workplace, for instance through distracting workers from their tasks unnecessarily often or at inopportune times. Table 1 shows a small selection of interruption research from the last few years, with a short explanation of what the point of the research is and a short explanation of the primary task used.

Table 1. A selection of recent papers on interruptions, with authors, year, the context of the research (i.e. what technology was used and/or what environment was used), and the focus point of the research, i.e. what effect the author(s) were examining.

Researcher(s), Year	**Primary task and context**	**Research focus**
Morgan and Patrick, 2013	Tower of Hanoi puzzle	Can the solving of a puzzle be made more resilient against interruptions by adding a slight delay to accessing the problem (increased access cost)
Baethge and Rigotti (2013)	Nurses, self report into a diary	Basic research into how interruptions affect irritation and perceived performance in a hospital environment
Cane, Caughard, and Weger (2012)	Reading a book, interrupted by spoken audio	Does a slight lag in responding to an interruption help in resuming the task
Jones, Gould, and Cox (2012)	Specialised software, desktop computer	Does a visual cue help with resuming an interrupted task
Leiva, Böhmer, Gehring, and Krüger (2012)	Mobile, app. Inferred from a provided dataset from a specialised app.	Measuring the effect of interruptions in natural context.
Arroyo and Selker (2011)	Desktop, multiple software running	Measuring the effect of using a disruption

		management framework when faced with multiple interruptions
Warnock, McGee-Lennon, and Brewster (2011)	A card-matching game of memory, using a notebook computer in a traditional office/lab setting	The effects of multi-modal interruptions
Altman and Trafton (2007)	Computer game, desktop computer	Building a model of how repeated interruptions affect performance
Gonzáles and Mark (2004)	Desktop computing, multiple tasks in multiple software	Observing how users manage multiple tasks and multiple task groups

An interesting point to be taken from table 1 is that artificial tasks bearing no relation to real world tasks still account for almost all research on interruptions. This is *not* a criticism of the use of artificial tasks, as these can be very useful for observing and modelling basic elements surrounding human cognition and behaviour, which is often necessary for basic research. This only shows that research on interruptions is in its infancy, and that the research needs to branch out into modelling modern contexts if it is to be of use in practice. This highlights the constant struggle between basic research and applied research, but in this case there is precious little practical research to be found. Moreover, the research that exists is in many cases contradictory, and in many cases does not show a clear a theoretical grounding. Some exceptions to this can be found, and frameworks exist that can be used as a basis for research. This includes the *goals-activation model (GAM)* (Altmann and Trafton, 2002), which has been widely used, as well as the *distraction conflict theory (DCT)* (Speier, Vessey, and Valaich, 2003). DCT suggests that different levels of task complexity create different levels of tolerable stress, with simple tasks allowing for higher levels of arousal before performance begins to drop. The goal activation model suggests that interruptions are dealt with through suspending the primary task after encoding the current state and rehearsing that while dealing with an interruption (Altman and Trafton, 2002).

When examining relevant and recent papers found on managing interruptions in a modern computing environment, such as Arroyo and Selker (2011), it becomes apparent that the lack of state of the art research creates problems for researchers, as most references in the papers introduced in table 1 are to papers written between 1999 and 2005. Arroyo and Selker (2011) are not thus alone in this, and those that make use of newer material, such as Leiva, Böhmer, and Kruger (2012), mostly reference papers on usability and user interface design, as opposed to papers on the effects of interruptions on people in modern computer usage contexts. Jones, Gould, and Cox (2012) include a number of more recent references, which shows that there is work being done, but even much of that work uses the kinds of artificial tasks that have already been discussed. Memory-game tasks, card sorting tasks, and highly artificial office tasks need to be complemented with real-world tasks as well as measurements and observations from real-world use of systems. This makes Leiva, Böhmer, and Kruger (2012) stand out, as the use of a large dataset sampled from mobile devices using an app that runs in the background and records user activity is a prime example of data about real behaviour. The problem for Leiva, Böhmer, and Kruger (2012) is that this was an externally sourced database, and the authors had no way of verifying quality of the information sampled. This is still an obvious step in the right direction, and suggests interesting methods of sampling data in the future. Of course, sampling massive datasets that show task performance does lack any information about affective states, or reasons behind changes in task performance. More research is required for that, with observational studies being one candidate for useful research that can be combined with data-driven research, even if that is a work intensive approach.

Older research exists that examines interruptions in general or in the workplace, but when seeing the word *workplace* being used it is most often possible to replace it with *in an office environment*. Noteworthy is also work that has been done on classifying interruption methods, such as McFarlane and Latorella (2002) who created a classification scheme for interruption coordination methods that take into account the requirements of both the primary task and the interruption task, and Spiekermann and Romanow (2008) who give a very thorough account of the state of interruption research up until that point. There exists good quality research on the effects of interruptions on people, but this research has largely been confined to older types of computer systems and has mostly used research methods that are artificial in nature or have a strong focus on the office environment. This kind of research is not likely to support the current generation of mobile devices, which support other usage models, and will be examined more closely in the following section.

CURRENT DEVICES AND INTERRUPTION METHODS

Having looked over existing research, it is also necessary to examine what the devices and interruption/notification systems that are currently in use are like. This is done in some detail here, and starts of with an overview of sensor technology and then goes into current interruption/notification system design.

An important change in the last decade has been the rise of mobile devices of all sorts, which all incorporate communications technologies and advanced sensors. Many of these technologies only became sufficiently miniaturised in the past decade or so, and were until that point difficult or impossible to incorporate into handheld devices. These communication technologies include, but are not limited to: Wi-Fi, Bluetooth (BT), Near Field Communications (NFC) for very close range communications, Infra-red, Bluetooth Low Energy (BTLE), and 3G and 4G cellular connections. The sensors incorporated include, but are not limited to: gyro sensors that sense orientation, three-dimensional magnetic flux sensors that work as a compass, accelerometers that sense motion of the device, light sensors, sound sensors, position sensing equipment such as GPS receivers, camera sensors, air pressure sensors, and more. It is important to understand that miniaturisation does not just mean that a sensor will be used for the same purpose as it was in the past, and that putting together a complex sensor array that can be made to work together also changes the possibilities. An application for taking panoramic photos is an example of this. This kind of application uses a mixture of relative orientation and motion of the device (gyros, accelerometers), absolute orientation (magnetic flux sensor), and location (GPS) to make seamless composite images that may even identify elements in the surrounding based on database information. This is a use that shows how miniaturisation can effectively change what can be done, as without having *all* these sensors working in unison then this type of use is simply impossible.

The current trend in the design of user interfaces for mobile devices or wearable computing devices is veering towards minimalism and clarity, and away from the more highly skeuomorphic designs that have dominated the last decades. This comes about at least partly because of the need for displaying information in otherwise crowded environments and with a limitation on the size of the displays in use. One result of this is that symbols such as icons are getting more simplified, even as the hardware capabilities of the devices increases. Modern phone displays can show detail above the limit of the human eye for distinguishing the pixels at typical viewing distances, yet instead of adding more graphic details the opposite approach is increasingly being taken, as can be seen with both the Windows Phone 8 system (Microsoft, 2013) and the iOS 7 system (Apple Inc., 2013a). As such, icon design is at its simplest when identifiable objects are to be represented; a simplified picture of that object can be displayed. As soon as more abstract concepts or cognitive actions are to be displayed the design of an understandable icon becomes more complex (Cooper et al., 2007). The current design approach seen in Windows Phone 8, iOS 7, Windows 8, Ubuntu, new additions to Mac OS 10.8, and more show that designers are aware of this.

Then there are other interfaces where simplicity is necessary because of hardware restrictions or cognitive issues. One example of this is the Google Glass system of wearable glasses that show a persistent heads up display that has to be made in such a way that it does not distract the wearer from his environment (Google, 2013a). The notification system used in Google Glass will be examined in more detail after examining general notifications in desktop and mobile operating systems.

Examining current operating systems gives clear indicators of how much attention has been put on improving notifications in recent years, how far notifications still have to go, and what computer users have come to expect from their systems. The most common notifications used are popup balloons of some sort, also commonly referred to as *toasts*. These toasts are commonly used by programs that give alerts as to their status, and can also have interactive elements. This solution is used in Windows XP, Windows Vista, Windows 7, Ubuntu, and other forms of

Linux in such a way that the toasts pop up or down from a global menubar that covers the width of the monitor. MacOS X (10.8 and onwards) uses toasts that appear without a link to a menubar, but give access to a unified notification centre that resides off the right hand edge of the screen (Apple Inc., 2013). Toasts are described by Microsoft (2013b) as:

> "A toast notification is a transient message to the user that contains relevant, time-sensitive information and provides quick access to related content in an app. It can appear whether you are in another app, the Start screen, the lock screen, or on the desktop. Toasts should be viewed as an invitation to return to your app to follow up on something of interest. Toast notifications are an optional part of the app experience and are intended to be raised only when your app is not the active foreground app."

Toasts are thus appropriate for showing notifications on large screens where the information presented in the toast does not obscure the task that the user is working on, but are more problematic when used on small mobile devices such as smartphones.

Operating system makers try to make their operating systems easy to use, which means that programs need to have similar behaviours and interaction methods. The organisations behind supporting the carious operating systems each have some form of guidelines that are freely available and aim to support software engineering on that platform. Apple and Microsoft both make extensive guiding documents both for their desktop operating systems and for their mobile operating systems (Microsoft, 2012, 2013; Apple Inc., 2013, 2013a) while Google has a less formal and less detailed approach for Android (Google, 2013a) and Google Now (Google, 2013) in the form of web pages, and Canonical makes available only programming guidelines and lacks any detailed usability design guidelines for the Ubuntu operating system. Microsoft provides guidelines as to how to use notifications, and suggestions in the Microsoft Windows UX guide show that the company's interface designers are acutely aware of the problems associated with notifications:

> "Don't abuse notifications:
> Use notifications only if you need to. When you display a notification, you are potentially interrupting users or even annoying them. Make sure that interruption is justified.
> Use notifications for non-critical events or situations that don't require immediate user action. For critical events or situations that require immediate user action, use an alternative UI element (such as a modal dialog box).
> Don't use notifications for feature advertisements! "
>
> (Microsoft Windows UX guide, 2010, pp. 13)

and:

> "Notifications inform users of events that are unrelated to the current user activity."
>
> (Microsoft Windows UX guide, 2010, pp. 35)

The official Android notification design guide is interesting in that it shows very graphically how a notification should be visually designed on an Android device and explain the technical side of how to display a notification, when to display a notification, and how different levels of priority can be set, but the interesting part is that no mention is made of what a notification actually does to a user. That notifications are an interruption mechanism that disturbs the user is not discussed except as a practical issue in the "When not to display a notification" section, and this sets the Android notification guidelines(Google, 2013a) in stark contrast to those from both Microsoft and Apple who both discuss extensively and explicitly that notifications are an interruption mechanism that interrupt the user's current task.

Apple, to our current knowledge, has the only desktop operating system that has a full notification system built in that integrates all aspects of notifications, yet even this system does not go much beyond pre-selecting whether a certain type of notification is permanently displayed in the centre of the screen (modal dialog), permanently or temporarily shown in the corner of the screen, or selecting which notifications can use sound to garner attention. In short none of the notification systems in use in general operating systems take into account even research from before the year 2000 (McFarlane, 1999) about user notification coordination, and certainly no attempt is made to have these notification systems interrupt only at appropriate moments. An example of this is that the notification system in Mac OS 10.9 does not take into consideration that a user is watching a film expanded to fill the screen (full-screen mode) and only display critical notifications, but rather pops any old notification onto the screen unless the user specifically disables the notification system. This is true for the *do not disturb* feature both on Apple's desktop operating system, OS X 10.9, and their current mobile offering, iOS 7.

Other human interface guidelines (HiG) for operating systems are thin on the ground. Ubuntu, the most popular Linux distribution, has no HiG document, nor does the Unity desktop environment used by Ubuntu. The desktop environments KDE and GNOME both have some human interface guidelines, but these are very limited when it comes to notifications. Gnome, for instance notes that "Using the status notification area applications can notify the user of non-critical events" and goes on to give a few examples of how to implement notifications for some tasks (The Gnome Project, 2012). KDE provides somewhat more detailed information that explains what each provided notification method should be used for (The Gnome Project, 2012). KDE similarly has limited information, with only minor technical listing of functionality provided (McBride, 2013). This is somewhat indicative of the problems with UI design for Linux, seeing as how the current KDE system (KDE 4.x) was released in 2008.

Notifications in current operating systems can thus currently all be considered extremely intrusive and "rude" in that they do not take the users' activity into account in any meaningful way, but rather all use what McFarlane and Latorella (2002) refer to as immediate interruption for all notifications, regardless of what the user is doing. This goes counter to the concept of considerate user interfaces (Cooper et al., 2007). This is what is going on at the operating system level, and the recommendation of each operating system maker is that the built-in notifications should be used for system-wide notification. Indeed, even if notifications are customised when using a program the operating system may still force the use of the its own notification system when an application is running in the background. This then needs to be taken into account.

CURRENT CUTTING EDGE IN THE CONSUMER SPACE

Special mention should be made of Google Now, which is a mobile, context aware, and smart notification system made by Google, is available on both Android and in limited form on iOS operating systems, and a version of which, called Live Cards, works on the Google Glass wearable system (Google, 2014). When comparing available notification systems to what can be found in research literature then Google Now looks to be the closest to the state-

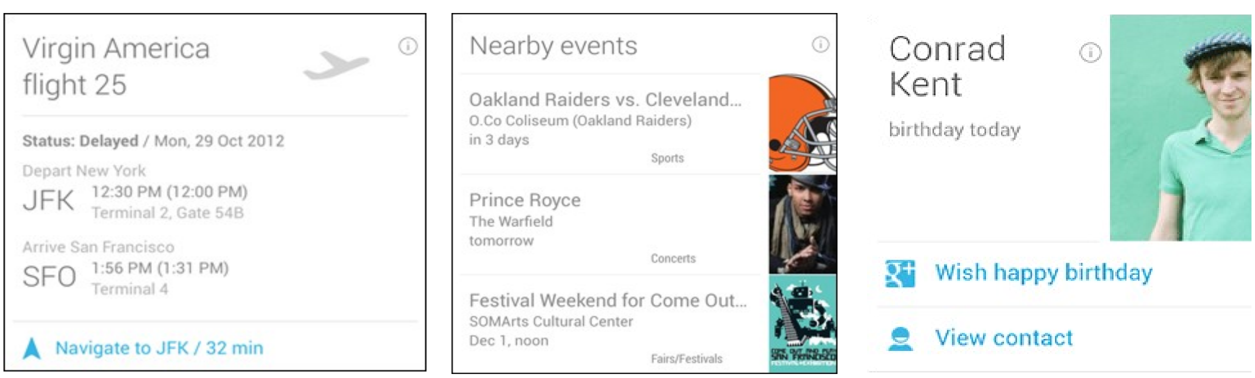

Figure 1: Examples of three different Google Now Cards. Flight information, a location based notification of nearby events, and a birthday notification. Note that all these cards offer direct interaction. Google 2013, Creative Commons Attribution 3.0 License

of-the-art that is currently available. Google now uses standard notification message templates called *Cards*, which allow applications to send information to these standardised cards for display to the user. Cards can be interactive, and there are a number of different Card templates that are formatted to support different classes of data, see figure 1. As a Card showing a birthday notification does not have the same information display needs as a card notifying the user of nearby events, these templates then allow for a simple API for developers that also support the needs of the user. As this is integrated with Google's various information services as well as using location services it is possible to implement interesting notification systems in Google Now. However, Google Now does not use notification suppression, aggregation, or any other tools for creating considerate user interfaces (Cooper et al., 2007) and is in this regard no more advanced than the other available notification systems. Google now is integrated with the notification centre in Android, as well as being integrated with the Google Glass system wherein the Cards are used to display notifications on the persistent heads up display offered by Google Glass. The Google Glass system is notable in having the potential to change how user interfaces in the general population as well as in industry can be used. In its current incarnation, Google Glass has a screen resolution of 640x360 pixels, and has a 3-axis gyroscopic sensor, accelerometers, a magnetometer (compass), proximity sensor, microphone (for voice commands and audio recording), Bluetooth, ambient light sensor, and audio output through bone conductance (so no in-ear device is required) (Google, 2014). Google Glass has recently been updated to incorporate wink (a long blink of the eye) controls for taking pictures (BBC, 2013), which shows an example of how eye movements can be used in a computer system. This combination of audio/video input and output, coupled with the advanced motion and environment sensing capabilities allows for exciting future development of head mounted information systems for industry. One such possibility would be to integrate a worker's (line worker or foreman) information flow, or even to show assembly instructions in a worker's field of view based on what custom assembly is currently required at a certain workstation. If such a system were to be developed it would require very careful design of the information management system, particularly any kind of interruptions such as notifications, or the system could become a liability. Any system designed for more current computing paradigms such as mobile handsets, desktop machines, or touch terminals would also benefit from such careful interruption design, and if well executed then an information system designed for these current devices is likely to support can easily be adapted to next generation computing systems such as head mounted displays.

It is clear from looking at the *current* state of technology that research on interruptions and research into the effects of the current technology in general has fallen far behind.

TECHNOLOGY TRENDS OF THE NEAR FUTURE

Mobile devices have already become ubiquitous, and can hardly proliferate much further, and their capabilities are likely to evolve during the next few years.

Wearable Computing

Wearable computing looks like an obvious group of new devices that must be mentioned. Notable among those are heads-up displays that can worn all day and can display information in users' fields of vision such as Google Glass (Google, 2014), and smart watches which can be used as computing devices or as satellite displays/interaction device that work in conjunction with another device which functions as a primary computing device (Pebble, 2014; Samsung, 2013). Less obvious is the use of smart textiles that can be integrated into clothing and can gather sensor data about the wearer, or even provide feedback through the use of light emitting fibres (Jansen, 2013).

New requirements

The new devices that either already exist or are on their way bring their share of new problems. Some of these problems may even be truly new and not just a repackaging of an old problem. Having dynamic information visible in the field of view at all times, for instance, is something that has never before been a possibility, except for limited amounts of time and in very specialised environments such as fighter aircraft cockpits. Yet the problems associated with this are at their root not technological problems, but rather problems of cognition. This means that examining existing basic research is necessary, but as this particular version of the problem may have special needs it becomes necessary to validate any existing research to make sure that it can be applied. The requirements for research created by current and coming devices are clear, and mostly unsupported. They include:

Research into the effects of interruptions of multiple classes and cognitive difficulty levels of primary tasks in a natural context. The effects that must be researched include task-switching times, error rates on multiple classes of primary and interruption tasks, changes in affective states, user acceptance, and ways of mitigating negative aspects.

Mitigating negative aspects is a cognitive and ergonomic issue that may need technical solutions, so a separate arm of research must develop *systems* for mitigating the cost of interruptions. This cannot be done effectively without having access to basic research into the effects of interruptions in a natural context, so the starting focus has to lie there. There is still nothing to stop developers from doing both concurrently, but that increases the risk of investing large amounts of time into system development that results in a far from optimal system.

Required research

It is hard to see any way that existing systems take human cognition into account when interrupting the user, and it likewise becomes clear that the research on interruptions is lagging behind the technical evolution that can be seen in both consumer electronics and specialised devices for various professional environments, such as in manufacturing. It is important to note that expanded research may also be required in the office domain, as there are massive changes there as well as in other domains. Although the lack of basic research in other domains means that it is required to first get the basic research from office to other domains, the introduction of all the same new classes of mobile devices into the office domain also creates new requirements for research in that domain.

When talking about all these domains that require research it is important to include diverse workplaces, such as manufacturing, but it is also important to examine general use outside of work, such as recreational use of tech using interruption mechanisms. How does that affect the enjoyment of recreational primary tasks? This is interesting from a quantitative perspective, i.e. does it shorten the time spent on recreational tasks? Does the user make more errors? It is also interesting from a qualitative or user experience perspective, which can ask questions such as: Do interruptions decrease the enjoyment taken from recreation? Do different types of interruptions or interruption modalities affect enjoyment? There are many possibilities here, and the field is young and expanding. Not to mention exciting!

Looking at McFarlane’s (2002) taxonomy of interruption coordination methods it becomes apparent that just to support existing models for interruption management there are a few areas that would benefit from research, as even these interruption management methods are not supported in any system that is generally available. To be able to provide any management of interruptions, the system must first be aware of the user’s requirements, and the user’s requirements vary based on what the user is doing and where the user is located. This suggests some basic required research that is required, and some aspects of which are listed here. The listed aspects consist of components that can be argued to be a part of context awareness.

Location awareness: Research is on-going for creating location aware devices and software. A major issue here has been creating the technologies required for a system to have high enough precision in the determination of location to be useful, both outdoors and indoors, as well as using low energy so that a device is useful for the appropriate length of time. Another factor is the range at which these location services should work.

A particular problem here has been determining location indoors, due to signal interference from walls and other indoors objects. Determining location at long ranges adds imprecision when there is any form of interference, such as indoors. A ten metre range indoors can thus be seen as a long range, while a hundred metres outdoors in an open field can be seen as a short or medium range.

The questions facing the UI and cognition part of research must therefore take into account some possible limitations and examine different levels of support for context aware interruptions based on location with regards to different levels of location precision. What this means is that different precision levels in determining a user’s location may give different possibilities for modifying notifications. An example of this could be an industrial workstation that can tell that the worker is 2 metres away from the workstation when it is appropriate to show a notification, and therefore enlarges the notification and minimises all other information on the screen. This presupposes a high precision of location awareness, at least in selected locations. There is also an ethical aspect to this, with personal information being sensitive and many people possibly wanting to avoid being tracked. This effectively means that any location determination that can be done on the user’s own device is more likely to gain acceptance than location services that rely on sending data to external systems.

Distance (proximity): Proximity can be treated as a part of location awareness, but can also have a separate component. Proximity can be determined without actually determining location, although if location has been determined then proximity has also been determined. This means that proximity can be a superset, with proximity services being usable to find whether a user is close to an item or far from it, which allows for customisations of UIs to support that. Uses for this are especially apparent for industrial application, such as workers in a manufacturing facility, where a UI could be customised to show appropriate information at appropriate scales based on worker proximity. The requirements for this are sensors for proximity, which can be done in a simpler manner than actual location sensing which needs a triangulation of sensors. Proximity sensing can also be done with low energy use.

Research into this is something that was not found in the literature, and requires both basic UI research and research into the effect of interruptions if used in this way, such as what the effect is of having large scale notifications appear on workstations.

Task awareness: One aspect in creating a system that minimises the impact of interruptions involves identifying the task that the user faces and the interruption task.

Activity awareness: An altogether more tricky proposition is the task of developing activity awareness. This entails understanding what the user is doing, based on sensor input. This might use gyro sensors, sound sensing, location sensing, and more.

Preference and requirement classification: Determining user requirements and user preferences is the subject of on-going research and is required for effective management of interruptions.

Environment awareness: Having the system be aware of the environment is linked to activity classification, but also quite separate. This entails having the system take into account requirements imposed by the environment, as determined by sensors (audio, video, motion) and location.

Discussion

The communications and information devices that are available today are without precedence; never before in the history of humankind has it been possible to communicate across vast distances wirelessly and receive those communications in multiple ways, almost wherever we are. What sets this form of communication apart is not just the power of the communications systems, but also that messages are sent without the sender having any understanding or feeling for the appropriateness of the timing of the message. This is new from a social and cognitive aspect as well as a technical one, and provides new and exciting possibilities for research from multiple perspectives. This short paper focuses on a limited subset of this, that is, only on what is missing in supporting the cognitive aspects of interruption management in information and communication systems.

Possibilities for research exist both for creating systems that will clearly be useful right away, as well as more abstract research with a less clear focus on current needs. That latter class of research is necessary to have a foundation for creating the next generations of communications systems, while the former are necessary just to make the current generation of communication systems less troublesome. An important factor to bear in mind when performing any such research is have a clear theoretical approach from the outset, and not reinventing the wheel every time.

The main suggestions from this examination of research into interruptions are that the current research needs to be on basic research to set up a theoretical foundation so that the coming generations of technology can be made to support human cognitive limitations. The effects of interruptions on human cognition are not going to change, as the limit is created by the human cognitive system. Research therefore needs to focus on basic human cognitive abilities, coupled with specific use cases, not on specific devices or device categories. If interruptions can be adequately managed in any pervasive computing paradigm, then the research backing that is likely to support other pervasive fields to a certain degree. This is in contrast to the current research that mostly focuses on limited work tasks with low risks associated with errors or failure, such as can be seen in the office environment.

The research has fallen behind the available technologies and the research suggested in this paper suggests some directions to catch up, and other directions to get ahead of technological advances. The scientific community must attempt to anticipate the needs of society, and are in many ways the best equipped to do so due to having access to the latest research.

Acknowledgement

The authors wish to thank the Sense&React Consortium for their input and their support. The Sense&React project is funded by the EU grant FP7-314350.

References

Altmann, E. M., Trafton, J. G. (August, 2004). Task interruption: Resumption lag and the role of cues. Proceedings of the 26th annual conference of the Cognitive Science Society

Altmann, E. M. (2013). Fine-grain episodic memory processes in cognitive control. *Zeitschrift für Psychologie, 221*, 23-32. (Special issue: task switching).

Altmann, E. M., Trafton, J. G. (2007). Timecourse of recovery from task interruption: Data and a model. *Psychonomic Bulletin & Review, 14*, 1079-1084.

Andrews, A., Ratwani, R.,, Trafton, G. (2009, October). Recovering from Interruptions: Does Alert Type Matter?. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* (Vol. 53, No. 4, pp. 409-413). Sage Publications.

Apple Inc. (2012) *OS X Human interface guidelines*. Accessible on the internet: https://developer.apple.com/library/mac/documentation/UserExperience/Conceptual/AppleHIGuidelines/OSXHIGuidelines.pdf [Accessed 2014.02.19].

Apple Inc. (2013) *iOS Human Interface Guidelines*. Accessible on the internet: https://developer.apple.com/library/ios/documentation/UserExperience/Conceptual/MobileHIG/MobileHIG.pdf [Accessed 2013.07.10].

Arroyo E., Selker T. (2011) Attention and intention goals can mediate disruption in human-computer interaction, in: Campos P., Graham N., Jorge J., Nunes N., Palanque P., Winckler M. (Eds.) *Human-Computer Interaction – INTERACT 2011, 13th IFIP TC 13 International Conference (Lisbon, Portugal, September 5-9, 2011): Proceedings, Part II,* Heidelberg: Springer, 454-470

Baethge, A., Rigotti, T., (2013) Interruptions to workflow: Their relationship with irritation and satisfaction with performance, and the mediating roles of time pressure and mental demands, *Work & Stress, 27* (1), 1-21.

BBC (2013) *Google Glass update lets users wink and take photos.* Accessible on the internet: http://www.bbc.co.uk/news/technology-25426052 [accessed:2013.12.30]

Bjelica, M. Z., Mrazovac, B., Papp, I., & Teslic, N. (2013). Context-Aware Platform with User Availability Estimation and Light-Based Announcements. *IEEE Transactions On Systems, Man, And Cybernetics: Systems, Vol. 43, No. 5*, September 2013.

Cane J. E., Cauchard F., Weger U. W. (2012) The time-course of recovery from interruption during reading: Eye movement evidence for the role of interruption lag and spatial memory, *The Quarterly Journal of Experimental Psychology, 65* (7), 1397-1413

Coraggio, L. (1990). *Deleterious effects of intermittent interruptions on the task performance of knowledge workers: A laboratory investigation*. Unpublished doctoral dissertation. University of Arizona.

Fokker (2012) *Fokker 50 Aircraft operating manual.* Eindhoven: Fokker.

Google (2013) Google Now: Actions in the inbox. Accessible on the internet: https://developers.google.com/gmail/actions/google-now [accessed 2014.02.16].

Google (2013a) Notifications-Android Developers. Accessible on the internet: http://developer.android.com/design/patterns/notifications.html [Accessed 2014.02.10]

Google (2014) Glass developers - Google developers: User Interface. Accessible on the internet: https://developers.google.com/glass/design/ui/index [accessed 2014.02.16].

Gnome (2012) Using the status notification area. Accessible on the internet: (https://developer.gnome.org/hig-book/stable/desktop-notification-area.html.en [accessed 2013.12.10]

Jansen, B. 2013. *Composing over time, temporal patterns – in Textile Design*. University of Borås, Studies in Artistic Research No 5, 2013.

Jones S. A., Gould S. J. J., Cox A. L. (2012) Snookered by an interruption? Use a cue, in: *People and Computers XXVI: Proceedings of the 26th Annual BCS Interaction Specialist Group Conference on People and Computers (BCS-HCI '12)*, Swinton: British Computer Society, 251-256.

Leiva L. A., Böhmer M., Gehring S., Krüger A. (2012) Back to the app: The costs of mobile application interruptions, in:

Proceedings of the 14th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '12), New York: ACM, 291-294

McBride, M. (2013) *System Notification Settings.* Accessible on the internet: http://docs.kde.org/stable/en/kde-runtime/kcontrol/kcmnotify/index.html [accessed 2014.02.18].

McFarlane, D. (2002). Comparison of four primary methods for coordinating the interruption of people in human-computer interaction. *Human-Computer Interaction, 17*(1), 63-139.

McFarlane, D. C., Latorella, K. A. (2002). The scope and importance of human interruption in human-computer interaction design. *Human-Computer Interaction, 17*(1), 1-61.

Microsoft (2012) *Windows 8 User experience guidelines*. Accessible on the internet: http://download.microsoft.com/download/e/1/9/e191fd8c-bce8-4dba-a9d5-2d4e3f3ec1d3/UXGuide.pdf [Accessed 2013.08.05].

Microsoft (2013) *Get started with Windows Phone 8*. Accessible on the Internet: http://www.windowsphone.com/en-US/how-to/wp8/start/get-started-with-windows-phone-8 [Accessed 2013.07.22].

Morgan P. L., Patrick J. (2013) Paying the price works: Increasing goal-state access cost improves problem solving and mitigates the effect of interruption, *The Quarterly Journal of Experimental Psychology, 66* (1), 160-178.

New Oxford American Dictionary 2nd edition (2005) New York: Oxford University Press.

Pebble (2014) *Discover Pebble.* Accessible on the internet: https://getpebble.com/discover [accessed 2014.02.19].

Samsung (2013) *Samsung Galaxy Gear jet black*. Accessible on the internet: http://www.samsung.com/us/mobile/wearable-tech/SM-V7000ZKAXAR [accessed 2014.02.19].

Schilit, B., Adams, N., & Want, R. (1994, December). Context-aware computing applications. In *Mobile Computing Systems and Applications, 1994. WMCSA 1994. First Workshop on* (pp. 85-90). IEEE.

Spiekermann, S.,, Romanow, A. (2008). Attention & Interruption Management for Systems Design–A Research Overview. Discussion paper.

Warnock D., McGee-Lennon M. R., Brewster S. (2011) The impact of unwanted multimodal notifications, in: *Proceedings of the 13th International Conference on Multimodal Interfaces (ICMI '11),* New York: ACM, 177-184.

Zeigarnik, B. (1927). On finished and unfinished tasks. *Psychologische Forschung, 9*, 1-85.