



The Future of the Mobile Internet: Lessons from Looking at Web Use

Abigail J. Sellen, Rachel Murphy
Information Infrastructure Laboratory
HP Laboratories Bristol
HPL-2002-230
August 30th, 2002*

E-mail: Abigail_Sellen@hp.com, Rachel_Murphy@hp.com

mobile devices,
internet,
appliances, web
use, user studies,
design

The Future of the Mobile Internet: Lessons from Looking at Web Use

Abigail J. Sellen & Rachel Murphy

Hewlett-Packard Research Labs

Filton Rd., Stoke Gifford

Bristol UK BS34 8QZ

Abigail.Sellen@hp.com, Rachel.Murphy@hp.com

Just what does it mean for the internet to “go mobile”? If we look around at some of the recent attempts to bring the internet to mobile devices, we might be forgiven for thinking it means a poor approximation of what we can do on our desktop Web browsers. For example, WAP phones give us access to the mobile internet, but the small screen and phone keypad input means that the capabilities of a “normal” Web browser are severely compromised. Many of the features we rely on for browsing are not there, displayed information is primarily text-based, and techniques for doing things like bookmarking are very different from what we might do on our non-mobile PCs. Combined with often shaky wireless network capabilities and limited WAP-enabled sites, we are presented with some of the features of Web browsing, but not all of them; with the ability to do some of the Web-based tasks we usually do, but not all of them. The result can be frustrating, especially in light of many of the claims that mobile operators made for the freedom that these new wireless internet services would provide us through the mobile phone.

When we look more closely at sales figures for WAP phones in the UK, they are quite impressive. A large proportion of people who buy a new mobile phone will buy a WAP-enabled phone. It seems that they do this not only because they want the most up-to-date phone, but also because the idea of internet access is appealing. After all, if surfing the Web through your desktop PC is useful, why would it not also be valuable on the move? Unfortunately the data also show that while many people own WAP phones, use of WAP-based services is low. One of the mobile operators reported recently that fewer than 5% of WAP phone owners actually use their internet-based services.

Contrast this with two of the success stories for the mobile internet: Research In Motion's BlackBerry and NTT DoCoMo's I-Mode. Both technologies have taken off: the BlackBerry mainly in the business market, and the I-Mode mainly in the teen market (and mainly so far in Japan). What is interesting about both devices is that they have been developed by companies with a clear notion of who their end users are, and which have designed their products and services around these users. In the case of the BlackBerry, both the device itself as well as its services have been targeted at business users, and at making the sending and receiving of email as easy as possible for business people. I-Mode has many more uses, but the approach here has been to evolve the design of a mobile phone into a technology that supports a cluster of new applications that teenagers might find compelling. A device with a screen larger than a normal mobile phone, DoCoMo capitalized on the fact that it was mainly Japanese teenagers who were flocking to IMode, and developed a set of specially designed applications with teens in mind. Today the most frequently used I-Mode services are those that are either entertaining, are social, or are both. Alongside its primary use for communication (via voice, email and voicemail), popular applications include on-line dating services, interactive games, and the ability to download ring tones, cartoons and images.

Looking at Blackberry and IMode, it is tempting to draw some early conclusions about the ways in which the mobile internet is most likely to take hold in the future. Technologies that have used viewed the mobile internet primarily as a mobile instantiation of what we already do on a desktop PC have met with little success. Technologies whose design reflects the use of the internet mainly as the enabling infrastructure for a set of applications aimed at a target user group are succeeding. In the latter case, this may call for alterations to the design of existing devices, for wholly new kinds of mobile devices, and for new interfaces and infrastructure to optimize the interaction with those applications.

Many of us would recognize this as the familiar “appliance argument”: that we need to streamline hardware and software for particular purposes by understanding and focusing on a well-defined set of user activities for target groups of people. Perhaps this is to jump the gun, however. It may be that more general Web browsing capabilities do have a place on mobile devices and it is simply the case that we need to work harder on optimising them for small screen interfaces. Perhaps it is the fact that the technology itself is simply not “there” yet and that improvements in wireless technologies making access faster and more ubiquitous will help make these services more popular.

An interesting issue, then, is to what extent we should rely on familiar notions of Web browsing when we think of the internet in a mobile context. After all, it may be that some kinds of Web activities are appropriate for mobile internet devices, while others are inherently unsuited. And it is also not enough to focus only on the suitability of such tasks. We need also to understand how different kinds of tasks might best be supported through mobile interfaces. Mobile devices have their own special constraints, small screen size and limited input mechanisms being the most obvious.

To answer this question, one thing we have found useful is to take a closer look at what people currently do through their desktop Web browsers. By looking at the *kinds* of Web-based tasks that people currently carry out, the *features* of those tasks, and the ways in which they fit into the larger context of other activities that people do in their day to day lives, we believe we can make more educated guesses about which kinds of Web tasks might translate well into mobile settings, and which, by their very nature, are activities that need to be carried out in desk-bound environments using more conventional computer interfaces. Likewise, we believe understanding such tasks more deeply helps us reason about what features mobile interfaces will need to have to support users more effectively. To do this, we also need to fall back on some understanding of what people do when they are mobile, so we can know something about what their needs and constraints might be.

How Knowledge Workers Use the Web

At HP, we’ve been doing research into both areas: into what people do with desktop Web browsers, and in another set of projects, into what people do when they are mobile (e.g., [1, 2]). Primarily our focus has been on work-related activities, although we have a keen interest in how domestic and social activities are woven into the working day. We have also been focusing up till now mainly on “knowledge workers”. Knowledge workers are people whose paid work involves significant time gathering, analyzing, creating or archiving information, information being anything from documents, policies, plans, and presentations to drawings, designs and graphics. Many of us in today’s working world are knowledge workers, and this segment of the working population is predicted to increase faster than any other in the future. Knowledge workers are interesting in that they tend not to work in a routine way. Rather, they show great diversity in what they do not only depending on their job, but also from day to day and project to project. To do this, they tend to draw on many different tools and technologies (some digital and others more mundane such as paper and whiteboards) to get the job at hand done.

One of our most recent studies was to look in depth at the kind of activities knowledge workers do when they use the Web (4). A primary motivation for this was that in looking at the published literature, there seemed to be a curious lack of any description of what it really means to “surf the net” or “use the Web”. It seemed to say little about the range of activities that people actually carry out with the Web, or about the defining characteristics of these activities. What are the most common or the most important Web activities people do? What other things do these Web-based tasks play into in the course of everyday life? There was simply not much written about these topics, and it seemed to us that without this understanding, it would be difficult to predict how Web use would change or evolve in a mobile context.

To do this study, we ended up recruiting twenty-four knowledge workers in a variety of different professional roles including: a magazine production editor, lecturer, building historian, information scientist, child psychologist, broadcast journalist, optical design engineer, and marketing manager. The study itself used a combination interview and diary approach that looked at the Web activities of these people over two consecutive days. Essentially this meant interviewing each knowledge worker twice, and on each occasion asking them to take us through all the Web activities they had carried out on the previous day (including any non-work activities). This was done mainly by looking at their Web history list, and sometimes by clicking through and revisiting the sites they had been looking at the day before. (Activities did not include email, since not all participants used the Web for email). By doing this, participants were able to tell us a “story” about each activity they carried out. This was captured in considerable detail: they were asked to describe each activity in terms of why they carried it out, what happened during the activity and any events leading up to or following on from it. We also recorded the amount of time they spent carrying out each activity.

A Taxonomy of Web Activities

When the study was completed, we had a detailed record of slightly less than 300 different Web activities (an average of about 6 per person per day). It was immediately clear that while there was much variety in the detail, these activities also fell into distinct clusters in terms of their goals and purposes, the way in which they were carried out, and the way in which they were done in the context of other, bigger and more complex activities. Overall, we found six categories of activity types:

- **Finding:** Using the Web to find something specific. Searching is goal-oriented and very well defined: E.g., Finding a fact such as a phone number, spelling or product name; a set of facts such as a list of ingredients for a recipe, or list of train times; or a virtual product or products such as a document, software, map, or image.
- **Information Gathering:** Less specific than “Finding”, but using the Web to purposefully research a specific topic for various reasons. E.g., Gathering information in order to compare, choose or decide about something (such as buying products or looking for jobs); in order to supplement a future task (such as collecting background information to write a document, or to prepare for a meeting); or in order to be inspired or get ideas.
- **Browsing:** Going to sites out of personal or work-related interest with no specific goal in mind but rather to be informed, stay up to date or be entertained. E.g., Browsing through a newspaper or magazine, following an interesting link, or checking to see what’s new on a hobby-related site.
- **Transacting:** Using the Web to execute a transaction securing future products or services: E.g., Making a bank transfer, paying a bill, ordering a physical product, or filling out questionnaires.

- **Communicating:** Using the Web in order to participate in chatrooms or discussion groups. (Email activities were excluded from this study.)
- **Housekeeping:** Using the Web primarily for checking or maintaining the accuracy and functionality of Web resources.

When we looked at the frequency of these different activities (Figure 1) and their average duration (Figure 2), one of the first things we saw was that Finding, Browsing and Information Gathering activities were much more frequent than Transacting, Communicating and Housekeeping. However, when we considered the most common activities, Finding and Browsing activities tended to be much shorter (just under eight minutes on average) compared to Information Gathering (on average taking 23 minutes). In terms of overall time spent on different activities, what we found was that these knowledge workers spent most of their time doing Information Gathering (Figure 3).




Figure 1. Frequency of different Web activities expressed as a percentage.




Figure 2. Average duration of different Web activities (minutes).




Figure 3. Overall time spent on different Web activities expressed as a percentage (frequency X duration).

A Closer Look at Different Web Activities

While these numbers are quite useful in giving us an overall picture of what people do, in many ways the detail is more interesting. For example, let's look more closely at three of the most common activities: Finding, Information Gathering and Browsing.

Finding

When carrying out Finding activities, our knowledge workers were quite clearly goal-oriented with very focused questions in mind. They asked questions like "What is the reference for this publication?" "What are the train times from Bath to London tomorrow?" "What is the curvature of this optic lens?" They looked for many different kinds of information. Most were *facts* such as names, dates, phone numbers, quotations, weather forecasts, procedures, prices, product names, financial data, schedules, and directions. A few were *images* such as maps, photographs and Webcam pictures. A few were instances of finding and downloading *virtual products* such as software, screensavers, images and documents. Most of the information they were after, however, ultimately took the form of single words, phrases, or lists of text.

The people in our study were quite experienced at this sort of searching, and, overall, quite successful at it. Perhaps too because of their experience, most of these sorts of activities were both short and self-contained. In other words, for almost all of the Finding activities we recorded, they were carried out without interruption and were done in one sitting.

Once they found information, an important issue was that of trust – whether they could trust the information they found to be up to date and accurate. They didn't trust "brochure sites" or sites perceived to be static and out of date. They

didn't trust sites where it appeared they were not talking direct to an organization but rather that the site had outsourced their Web services. In general, they didn't trust sites where they were uncertain of who produced the information and whether it was from a reputable source. Discussion often centered on the importance of using branded sites such as the Encyclopedia Britannica, or sites associated with official societies or universities. Such resources were usually bookmarked and used more frequently than others. Perhaps it is no surprise then, that these knowledge workers used general search engines only 19% of the time during Finding activities. Much more common was the use of bookmarked or familiar sites specialized around particular kinds of knowledge. This has important implications for the mobile internet, which we shall come back to.

A final important point about Finding is that it would be wrong to view these activities as meaningful tasks in and of themselves. Finding almost always led to further activities. For example, once information was found, people would then go on to do something with that information: If they found a phone number they would ring it; if they found some software they would download it. Some of the things they found they would email to share with others; some things they found such as images, references and quotations they would then incorporate into a document; and some of the items they found they would then go on to order and buy. Finally, a very common practice was to seek out information such as maps, procedures, and lists that needed to be kept, usually only temporarily, so that they could be referred back to during some specific time period, such as a journey. Again, in terms of design, it is useful to consider these follow-on activities, as we will come to later.

Information Gathering

Information Gathering, it turned out, was very different in nature, involving the collecting of information around a particular topic. This differed from Finding in that many of these activities involved a *set* of questions to be answered (e.g., "Who are the key players in this market? What other products do they offer? Where are these people located?"), or in that the questions were sometimes quite open ended (e.g., "What job opportunities are out there?"). Such activities often involved comparing and contrasting information across sites or organizations (e.g., "How does my insurance policy compare with others?"). In many cases, however, although these knowledge workers could sometimes specify a high level question they were posing, often they could not. Rather, they said they wanted general information about a topic (a country, a product, a company, a person and so on).

Interestingly, as with Finding, Information Gathering tasks usually began with bookmarked specialist sites, with general search engines used only 29% of the time. Here again, the use of trusted sites was a key issue. As an information scientist we interviewed put it, one needs to know about the *quality* of information--certain sites provide better quality information than others. Part of this knowledge worker's job was to know and assess which sites provided the best quality information and which did not. An important phase of his work was when he would engage in "scoping searches": time dedicated to evaluating sources, saving the valuable ones to folders. Later, he would revisit these to focus in on the documents buried within them. Other knowledge workers too, with quite different jobs, talked of the need to gather information from trusted, reputable sources. These would be bookmarked and used more frequently than others. When information came from an unknown site, it would be checked against other sources such as reference books. Finding good websites to visit was something they learned about from their peers. Cross-checking information also was something they would talk to colleagues about, in the workplace, or by picking up the phone.

While the accuracy of information was sometimes a significant issue in using the Web to gather information, this was not always the case. Some Information Gathering activities were for the purpose of inspiring ideas (for example journalists who were looking to see what the hot stories were and following up leads.) Others were simply to provide background information on a topic area. As a marketing consultant put it, the Web could be very useful in “building a picture” of an organization, topic or person.

Whatever the purpose, Information Gathering activities, unlike Finding, could be quite time-consuming and complex. They were essentially meta-level tasks usually involving navigating multiple links and sites. It was clear that this process was often very reliant not only on complicated navigational patterns, but also on scanning and skim-reading large amounts of material to assess its relevance. Another characteristic of these activities was that they often unfolded over time, sometimes spread out over days or even weeks. In fact, we found almost half of these activities were not completed in a single session, being interrupted either due to other demands or because of the amount of time required to complete them.

Because of this, the people in our study used various ways of saving the interim results of their Web activities. Some used virtual folders (either on the desktop or bookmark folders) to collect together links associated with different topics or projects. Others printed out their findings, including search results, tables of contents and home pages. Printing was especially useful when they were worried that interim search results would not be preserved or that information would change in the future, making it hard to get back to where they had been.

Browsing

Browsing differed again from both Finding and Information Gathering in that it was not generally goal-driven. Rather, the people we interviewed used the Web to see what was new, to keep up to date, or be entertained. This might involve information in a work-related field, from a home country, or in hobby-related sites such as science fiction, gardening and football sites. Over half of these activities were routine in that they were carried out either daily or weekly. In these cases, favorite sites either were bookmarked, entered from memory, or reached through automated email lists. The remaining activities were one-off browsing events, either as the result of emailed links from friends and colleagues, or through serendipitously discovering and following interesting links while at other websites.

Unlike Information Gathering (but similar to Finding), Browsing activities tended to be both short and self-contained in that they were hardly ever interrupted or left “mid-stream”. Information was also rarely printed out or saved in a virtual folder, and there were usually no follow-on activities except occasionally to bookmark or email an interesting item to someone else. In being activities “unto themselves” rather than leading on to other activities, Browsing was different also from Finding.

The kind of information browsed was here, as in all other categories, primarily visual, sometimes involving reading through pages of text and graphics, but usually in a superficial way. However, many of these items were of a smaller format such as photos, the weather, sports scores, jokes, and stock prices. People also heavily relied on headlines and summaries to look for items of interest.

A final important point is that, unlike the other categories, Browsing activities were the most likely to involve multimedia. About 10% of the Browsing activities we looked at involved audio, with a few cases involving video. Many of the participants said they wanted to make more use of multimedia, but found that they either had problems with the software,

or weren't willing to wait for the content to download. Only one person in our study actually used the Web to stream media (in this case, the radio) on the days that we interviewed them about.

Implications for the Mobile Internet

Perhaps the most obvious point this study makes is the extent to which terms like “surfing” or “browsing” the Web gloss over and even misrepresent what people use the Web for in a deskbound environment. There were many different reasons why the people we interviewed used the Web and many different behaviors they engaged in: some goal-oriented and some not; some centered on focused questions and some not; some short and self-contained and some spread over time. Not surprisingly, different activities involve different patterns of use and types of information. In doing so, they may be part of larger tasks, and they may draw on other resources such as documents, technologies, and even other people.

At the very least, this should cause us to be somewhat circumspect when claims are made that mobile technologies will allow us to surf the Web. Such claims can indicate that there has been little thought given to what *kind* of Web use they mean, and for what sorts of tasks. If they are not thinking about particular kinds of activities, we might well wonder whether they have thought about their usefulness in a mobile context, or whether they have thought about how best to support them.

What Activities are Appropriate on Mobile Devices?

For example, our research clearly shows that only some of the Web activities people currently do with deskbound Web browsers are suited to small, portable, wireless technologies. For example, many of Finding activities we saw could be easily carried out through a small screen. Much of the information the people in our study sought was of a small size and format, such as words, phrases, lists, or images. In addition, many of the kinds of information our participants sought (e.g., train schedules, phone numbers, directions and movie times) are those that people on the move might well value to help them get around, and to help them coordinate and plan events in their daily lives.

It is also clear that some of the Browsing activities we saw would also be appropriate on mobile devices. Like Finding, some of the information that users browsed through could be displayed with no difficulty on a small screen (e.g., jokes, still images, Web cam shots, sports scores, weather and the like). Here too, some kinds of video or animation might be suited to small screen displays. Audio is another clear choice for mode of delivery through mobile devices, and a kind of information that fits well with the things that people did while Browsing. We also know from our research on mobility that people on the move often have long periods of what they sometimes refer to as “dead time”—time stuck in queues, on buses, waiting in offices, train platforms, and other places where there is little to do (e.g. see 2). Entertaining or distracting services tailored for such devices, and designed for the kinds of things people do when Browsing, would be valuable in these kinds of mobile contexts.

The research also shows, however, that other kinds of activities would be very poorly suited to use through these devices. For example, some of the Browsing activities we observed were heavily dependent on the need to scan large regions of text and graphics, such as the time people spent scanning virtual newspapers, magazines and catalogs. It is quite difficult to see how this could ever be done effectively without a large screen display and without drastically altering the nature of the task. Likewise, when it comes to Information Gathering, the results show that such activities would be, for the most part, entirely unsuitable for supporting on small screen devices in mobile settings. These sorts of activities too often involved

scanning of large sections of text. They also often involved complicated navigation in and across sites. This in turn resulted in a need to manage and archive information from multiple sites. It is unlikely that any small screen device could successfully support this kind of complexity. The results also suggest that such activities may be hampered even on large screen devices in settings away from the desk (i.e., laptops) because of people's reliance on other resources and infrastructure to support the kind of meta-task activity that goes on in Information Gathering. In other words, the people we interviewed often had checked information against other sources (such as books, paper documents, and the knowledge of colleagues) and used paper folders, printouts, and notebooks to preserve and collate information in the face of tasks that could be both protracted and interrupted. Better on-line Web tools might be developed to help address these problems, but it is difficult to see how they could easily do so on the move, far away from the desks, books, documents, files, printers, and people that are such an integral part of working life. Such tasks seem fundamentally the sort that need to be done in fixed, predictable environments where there is a familiar working infrastructure in place and readily to hand, and where there is room to spread out information and artifacts to use them to think, plan and organize (for additional evidence on this point, see [3]).

Optimizing the Mobile Interface

By looking closely at only three of the Web activity categories we identified, we can start to more clearly focus on the kinds of things that people might successfully do in the world of the mobile internet, and that they may *want* to do. But there still remains the issue that to translate Web tasks from the context of a desktop Web browser to a mobile Web browser, one of the main challenges will be getting the design right for different kinds of mobile interfaces. Once we understand the task we are trying to support, however, it is easier to see how the design of a particular mobile service and device might be optimized. This may well mean substantially altering what we think of as an interface to a Web browser.

Consider bookmarking through a desktop Web browser. One of the things we have seen is that when people do Finding activities, they want to get quickly to familiar sites, they don't necessarily want to manage or archive Web information as they go in the way that they would do while Information Gathering. Conventional bookmarking techniques are designed to let users do both: to quickly get to a site they have been to before, and to save and gather together sites within hierarchical folders. Ironically, doing one thing works against doing the other: for example, collating too many sites in a folder means it takes longer to get to the one you want. A method of bookmarking designed for speed, not archiving would seem a better design approach to support Finding through a mobile device.

One concept we have been playing with has been the idea of mapping one's favourite sites onto the beveled dial of an internet-connected watch (see Figure 4). Users pre-select the most frequent sites they visit when wanting to Find or Browse. Because of the small display size, this would make sense only when the sought-for information is suited to presentation in a small format (the weather in London, a view of surf conditions on the south coast, a selected image from your database of favourite photos, your company's stock price).



(a)



(b)



(c)

The “InfoWatch” showing 3 possible favourite Web sites the user might want to access. Shown here include: (a) a site with the local weather; (b) a home-based Web cam image; and (c) keeping an eye out for the best deals on flight prices.

But we can also do much more to optimize user’s interaction with the Web through existing mobile devices. Consider the internet-enabled mobile phone. We have seen that when users do Finding activities on their desktop PCs, there is a set of follow-on things that users often want to do to after information is found. This suggests that there may be a very useful small set of functions that users might find valuable to get easy access to on their mobile phone once the desired information is located. For example, we saw that people often print out information they find such as directions, maps, lists, and procedures so that they can refer back to these during a shopping trip or a journey. This suggests that a single button facility to display a temporary store of information clipped from the Web might be a valuable tool. Users could save found information and instantly access it without having to connect to the Web again, and without cumbersome navigation through menus. This could then be easily deleted when it is no longer needed. As another example, the fact that phone numbers were often sought from Web sites and calls made in conjunction with Web activity points to the value of: instantly highlighting phone numbers on Web sites; allowing one-button calling from sites from a Web-enabled phone; and supporting quick storage of phone numbers from Web pages. Another example: The need to incorporate found information into documents suggests that another useful feature would be a networked clipboard into which references, quotations and images saved for the purpose of documents could be kept until the user was in a position to work on the document (e.g., at a desktop machine). Another function, suggested by the data, that might be useful especially on portable device would be the ability to “watch” or monitor any Web site or selected portion of a Web page for change. There are currently services which will do this for stock prices, but this could be a much more general kind of feature, for example allowing users to monitor airline prices, changes in Web-cam images, or indeed any update to a site.

CONCLUSION

We end by considering what our research has to say about the notion of a general browser. By considering three of the most common kinds of Web-based activities that people do, we have seen that different kinds of Web activities point toward certain kinds of devices, services, applications and infrastructure that best serve those activities. These can be at odds with each other: for example some activities are better supported on small-screen mobile devices while others are more suited to large screen devices in fixed environments rich with supporting infrastructure.

We need not stop there, however. Taking this to its logical conclusion fits well with the “appliance vision”: that new devices, applications and services specialized to particular classes of activity have the potential to provide interesting, effective ways of interacting with the Web. For instance, we have seen that when people browse (in the more accurate sense of the word), they browse through many different kinds of information. As a user, I may want to select the tools best suited to maximize my interaction with that information. I might want to do this by specifying and streaming different kinds of Web content to different kinds of appliances. For example:

- on-line magazines and newspapers go direct to my printer or e-book;
- shopping sites go to my portable home tablet;
- football scores, stock prices and maps go to my handheld or wearable computer;
- photos and Web-cam pictures go to the wall-mounted display in my kitchen;
- MP3 files go to my internet-enabled audio player;
- video goes to my television or seat-back in the car.

There are many advantages to this kind of approach. Because appliances are specialized and pre-configured for particular kinds of information, we can avoid long download times for high bandwidth content and the time to acquire the tools to deal with content (something the people in our study complained about). More importantly, I access the information where I am most likely to want it, and interact with it in a way best suited to the way I want to use it. Some of this will occur in mobile settings, and some not. When this begins to happen, what we think of as the internet, mobile or not, will no longer be shaped and determined by what we do through Web browsers. Chances are we will think of the internet less, and think more about the things that it enables us to do.

REFERENCES

1. O'Hara, K, Perry, M, Sellen, A and Brown, BAT (2001). Managing information on the move. In Brown, B, Green, N and Harper, R (Eds.) *Wireless world: Social and interactional aspects of the mobile age*. The Hague, Netherlands: Springer Verlag.
2. Perry, M., O'Hara, K., Sellen, A, Brown, B., & Harper, R (2001). Dealing with Mobility: Understanding access anytime, anywhere. *ACM Transactions in Computer-Human Interaction, Vol. 8, Issue 4 (Dec. 2001)*. ACM Press: New York., pp. 323-347.
3. Sellen, A.J. and Harper, R.H.R. (2002). *The Myth of the Paperless Office*. Cambridge, MA.: MIT Press.
4. Sellen, A.J., Murphy, R. M., & Shaw, K. (2002). How knowledge workers use the Web. *Proceedings of CHI 2002*, Minneapolis, MN. New York: ACM Press

