LONG PAPER

Ubiquitous search for smart workspaces

Daniel M. Russell

Abstract One key goal of personal information management is to let users search all of their information assets easily. As knowledge workspaces evolve, a need to provide ubiquitous search services that allow the user to search anything, from anywhere, at any time becomes increasingly important. A key feature of the smart workspace is the ability to support this access to personal content. Here, several current systems that provide some of these capabilities are examined, and implications for future information access systems are derived.

1 Introduction

Information access is central to nearly every part of information-based work environments. Access implies being able to find information to work on it, and the ubiquitous, distributed, intelligent workspace needs ubiquitous search.

That is to say, the Streitzian perspective on smart work environments implies an augmentation not only in distributed working spaces and content interoperation [24] but also in the ability to locate and find relevant information. Saving and organizing all information is one thing, browsing that stored information is nice, but being able to search across all relevant repositories and find when needed is critically necessary for successful use whenever information storage becomes even slightly larger than what short-term human memory can handle [7].

This paper argues that *ubiquitous search over the contents of the smart workspace* is now necessary and

D. M. Russell (✉)
Google, Inc., Mountain View, CA, USA
e-mail: DRUSSELL@GOOGLE.COM

expected component of a successful intelligent working environment. Information in use in the workspace needs to be indexed, searchable and also robustly available through many venues in order to support the highly distributed, interactive and highly information accessible models that are growing common in the workplace.

1.1 Previous perspectives on search-in-work

As Malone and others have pointed out, a physical filing structure is a hierarchical category structure [14, 18]. However, placing a virtual object into a tree structure is exquisitely sensitive to choices made while descending the hierarchy. An error made in choosing the category can result in an object being very far from its "correct" file location. And since filing structures frequently change in response to shifting tasks, organizations and task requirements, the probability of perfect filing grows ever smaller with time and changes.

Hierarchical filing systems do have a point—they define plausible scopes for organizing content. That is, content objects that are superficially similar can be organized within a hierarchy, with the hierarchy defining metadata properties. With a competent search tool that can search for objects within the entire computer, the user can always search the entire space. But when even a little bit of metadata is known (especially if encoded within the structure of the hierarchy), search can often be scoped to a subportion of the hierarchy, and therefore to a particular set of metadata values.

Personal and work-set information systems have commonly included search tools as part of their base function set. Calendars without search are still useful, but calendars *with* a search capability are often used in very different ways. (Try to find a particular appointment that occurred

sometime in the past 5 years of calendar data without search!)

Interestingly, Ducheneaut and Bellotti found that sorting e-mail as a way to find a particular message was a common—and perhaps dominant over standard search in Microsoft Outlook [4, 11]. However, as Google has found with Gmail, when search is fast and simple, usage of the search function to locate an individual message skyrockets. Whenever a document store contains significant personal information—be it e-mail, contact lists, a file store or physical files—search is a critical function. Thus, the need for search is fundamental for managing and using a large personal store. And when the data become large or sophisticated, the need for search tools becomes ever more important [12].

1.2 Basic issues, problems and challenges

Workspace information comes in many forms and distributed over a variety of working spaces, contexts and virtual locations. Workspace information has often been thought of as extensions of past practices—the physical desktop phone number files for contact information, calendars for scheduling, the handheld index card for notes and so on. In more current conceptions of work practice (e.g., [22, 23]), people are seen as working in complex landscapes of information that is ubiquitous, shared and distributed. The challenge is how to work within these new and constantly evolving frameworks.

Personal information required to support work practices comes in a variety of forms: e-mail, documents, contact management, personal journals and day logs, meeting notes, outlines, task management (from formal to informal to-do lists), lists of books to buy and groceries to pick up, class and soccer schedules, calendars of events by organization and family boundaries.

It becomes clear through scenario analysis of work practices that "personal information" comes in many kinds and forms, from different sources, and in unexpected ways. Personal information can be thought of as the *entire* collection of information that someone has within their immediate sphere of awareness. That is, personal information is more than just the contents of specialized databases or particular applications, but is effectively everything and anything that's within a user's controlled content space. Roughly speaking, everything that's on their "personal" computer, be that laptop or workstation—at home, at school, at work… or more likely, some combination of all. And, increasingly, personal information resides on several devices—PDAs, cell phones, multiple desktops—at the same time.

Personal information is also created by a user's behavior patterns both explicitly and implicitly. People often maintain bookmark lists of frequently visited web sites (another kind of personal information), but the bookmark list is a sample of their browsing and search history. This too is a kind of personal information—the history and paths of behavior that are collected implicitly and stored locally (e.g., by search browser history).

A basic question for personal information system design is understanding not just what personal information is, but how people would like to use all kinds of personally created, derived and observed information. What kind of information needs to be stored and how will it be used? What needs to be retrieved? In short, what are information hopes, needs and desires of users?

It is clear that people have a very difficult time anticipating future need or retrieval requirements for personal data. The Rolodex is a well-known system for storing phone numbers and addresses. Alas, it just does not scale well to many thousands of e-mail addresses (plus IM handles, plus mailing lists, plus passwords, etc.). Unfortunately, just as the Rolodex model does not scale well without augmentation, the increasing amount and varying kinds of personal information implies that future personal information will come in an increasing number of forms and ever-increasing quantities [9]. Personal information users need a capability that is robust and fairly general to accommodate future needs.

All in all, the challenges for personal information continue to grow and are being driven by new means of communicating, new devices and the shift in the definition of what constitutes personal information. As Landsdale pointed out long ago [16,17], the variety of personal information uses seen in real behavior is not well reflected by the available tools. User behavior is far richer than the tools they typically use. This disparity between what is needed, what is possible and what is available for use drives research in this area.

2 Work on workplace information access

2.1 What kind of work?

Workspace information has a long history of capturing and organizing information to make personal and work lives simpler. Searching within personal information has an equally long history [5, 15, 18, 19, 26]. As various trade-offs between capturing, storing, cataloging and indexing personal information have been explored, one thing stands out: There have always been a diversity of ways to store personal information. Clearly, after a non-trivial amount of information has been created and stored, getting back to it—searching it—becomes the next great task.

People have created information management tools to capture many special purpose kinds of information—contact information, personal health records, telephone

numbers, login accounts and passwords, bank information, to-do lists, etc. The list is as long as every special purpose information collection that people have created. Someone, somewhere, at some time has written a program to record and find that information kind.

Workers also commonly appropriate whatever systems and mechanisms are at hand to capture and make personal information findable. People often use e-mail as a mechanism to capture their personal information and, by extension, their workspace information [11, 27]. Why do people use this apparently non-optimal behavior? Because e-mail is ubiquitous, easily available, and most e-mail systems have a reasonable search facility that lets the user search, filter and sort by different properties. In addition, e-mail has often been a repository of content, since e-mail attachments permit the user to associate arbitrary text with a file (e.g., a picture or an otherwise unfindable blob). Shifting from a time when storage was expensive to an age of large amounts of essentially inexpensive storage, the cost of storing an item is no longer the dominant factor in determining storage vs. use, the cost of finding something again is. E-mail, as the great common denominator, has become the personal information manager for many [8].

Another focus of work has been to broaden the range of personal material that is captured and stored as part of the desktop space, making a very rich and robust personal content collection. Microsoft's "Stuff I've Seen" project aims to cast a broad net over the digital information space that people inhabit and provide rich new ways to capture even more information (such as automatically capturing images and integrating multiple personal data trails into a single personal data space) [12].

Personal information management breaks down to capturing information and then finding it again. Unfortunately, essentially every personal information tool's search system is particular; search has been idiosyncratic and specialized. While carefully tuned and specialized search mechanisms can be very effective when operating over controlled and moderated content domain (for example, a system like Lexis-Nexis, which is tuned for search in the domain of legal records), personal information by its nature often has unanticipated needs and unanticipated structure. Thus, search becomes a driver for personal information management. I find it difficult to accurately anticipate what kinds of information needs people will have and with the rapid growth of new media forms, of the primary problems of personal information management is that of search. How can search work in such a volatile environment?

2.2 A change in search interfaces

Search interfaces have always been specialized to their particular information content and style, featuring special access functions for dates and calendar operations, or for sorting along different properties and a variety of display formats. In order to support working styles and practices, the tools of personal information management were tuned for the tasks identified and to be efficient implementations for access and use. Searchable databases were primarily in the hands of large mainframe owners with carefully controlled content and index creation. With an increase in personal computing power, the stage was set for a shift in the way people thought about and used search capabilities.

Several years ago, a major change in search interfaces began with the introduction of web-based and desktop-based full-text indices and a fast, simpler way of making queries over a large unstructured corpus [19]. While keyword-based search had been around for many years (as embodied in systems such as DIALOG's pay-for-search system of newspaper and bibliographic data), less-structured text-based search quickly gained ground as the dominant model of search. This was never more true than when Excite, InfoSeek, Lycos and Yahoo launched free web search services in 1994 (with AltaVista following in late 1994). Very quickly, an expanding set of people began to understand search not as the construction of Boolean expressions over a rigorously defined database schema, but as a question of choosing the right "keywords" that would give back links to documents of specific interest.

The new, easily available capability had a profound influence on research in search. Until then, search front ends focused primarily on getting the expression builders to more accurately reflect the semantics of what the user intended. The biggest problem is that common natural language usage of AND and OR is not quite the same as their Boolean counterparts [1]. With free text as the predominant query model, tools other than web search began to adopt the new style for its simplicity and breadth of understanding by users.

Equally importantly, it is now understood that search of personal information stores has become a commonplace—a user could start to see how personal, local search would become a reality. The big question became how should one organize and index personal information? Several solutions have been proposed over the past several years, and we consider those in the next section.

3 Search as ubiquitous tool for work practice

3.1 "Search all content" as part of ordinary search

Just "search everything" is one answer to personal information management questions. For example, the basic Google search model is being brought into daily information-seeking behaviors. We know that for many users, web

search has become a common practice; it is the way people find the local pizza parlor, the latest news items or technical articles on arcane topics in computer science. And as personal information becomes increasingly integrated into the fabric of our personal data stores, searching everything becomes a popular information management strategy [7, 8].

Some desktop search systems have been built in the past to locate information distributed over a local store/local desktop application space. Systems such as the ones reported in [2, 10, 28] and the Google Desktop System made personal information search just an ordinary part of every "web" search done on a browser [13].

Normally, web search is just that searches applied to the contents of the World Wide Web. But when Google desktop search (GDS) is activated (by downloading a local client indexing application and letting it run in the background), the results of any Google search can have "local desktop" results blended into the results. In Fig. 1, the two best local desktop search results are shown as the first hits. To see the rest of the desktop results, the user would click on the "831 results" link. Through this approach, GDS makes the working contents of a personal computer into one large, searchable content repository that is effectively merged with the external web. Thus, when a user does an ordinary web search, they also simultaneously search their local desktop files. Since GDS indexes many types of files (e-mail, chat histories, web search history, Microsoft Office documents, PDFs, calendar appointments, etc.), the single action of searching (using a web browser) effectively searches the entire personal information space of the user as well.

In operation, GDS sends the users' query to two locations. One copy is sent to Google.com and performs a standard Google Web Search. A duplicate query goes to the GDS application running locally, which searches the local index. GDS then intercepts the results page before display, merging in local desktop results just above the web search results, so the user can see both at once.

The interface challenge is to give the user a sense of the locally stored information without displacing too many of the organic web search results. GDS includes a brief summary of the possible desktop hits, showing the total number of local results, along with two of the highest ranked local results. Each high-ranking result is displayed as a regular web-style link (blue and underlined), along with a summary snippet to the right. Unlike a standard web result snippet, the GDS snippet is limited to a single line of text, stretching from the end of the title to the end of the display region. This is a remarkably small amount of space in which to summarize a potentially large number of local results, but the trade-off is clear: The user's original intent (determined by their action) was to do a web search—GDS is only giving an indication of possibly relevant local results that the user might not have recalled.

GDS works by full-text indexing the content of files on the local repository, typically the "desktop." As is standard practice for such personal system crawler/indexers, GDS runs in the background, limiting its overall performance so as to not disrupt the user when doing their normal work. A full-text index is made of the file system contents (including transient objects such as web searches and chats), taking up around only a small fraction of the total file system objects being indexed.

Fig. 1 Google search of local personal information is invoked by doing a *regular* Google search. When desktop search is active, the desktop is indexed and the high-quality matching results mixed into the regular, organic, search results. Here, the top two best desktop results are shown next to the multicolored GDS "swirl" icon, with all 831 results available, including e-mails, notes, calendar entries and image texts
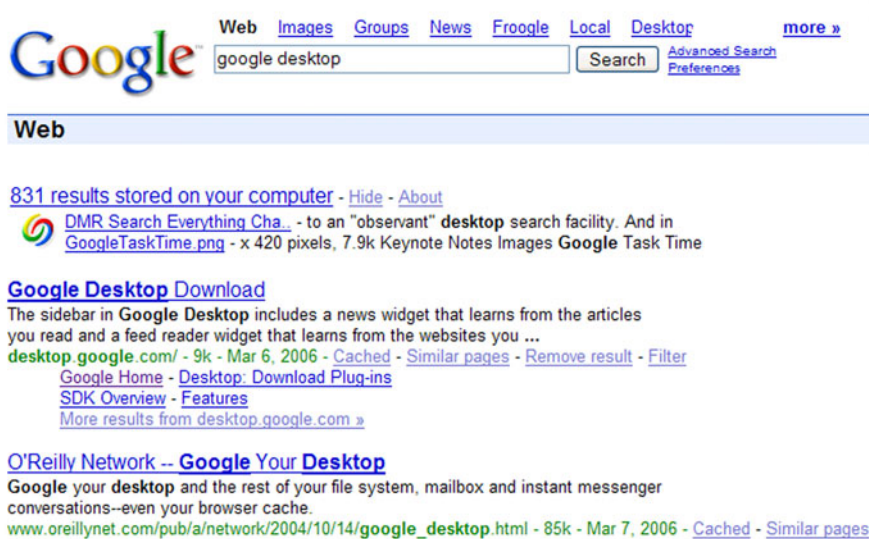
Crawling and indexing are fairly off-the-shelf technologies, with open-source implementations widely available [20]. The key value of GDS for personal information management is its integration with common activities (ordinary web search) and the various extensions that support a range of personal information tasks. In its current implementation, GDS is the key part of a larger package (named "Google Desktop") that also provides a number of other capabilities to manage personal information–multiple desktop indexing (a search run on one computer will automatically also search all "linked" computers with a shared index, making it possible to have a distributed working environment), a web sidebar for streaming information from multiple sources to the desktop, a scratchpad mechanism for taking fast notes and a set of accelerators to help share information with friends via IM or e-mail.

3.2 Scoping and broadening

Scoping means limiting a search to a particular range of things to be searched—all PDF files, all e-mails from Derek or just the files within a subdirectory. While this sounds like a bad idea, it is actually a useful way of looking for things that are common. For instance, searching for [June 24] over your entire personal information store is a terrible idea in general, but within the scope of the 2006 personal calendar, it can be a very effective way of getting to an item that the user is aware of. Not only would such a scoped search quickly navigate to that particular day, but it would also show up other references to that day that are hidden elsewhere in the calendar (such as a note written on May 2nd: "final paper due on June 24, 2006").

Most personal search tools implicitly scope their searches to the kind of information they are operating over. Each kind of search—the calendar search tool, a personal contacts search and a search of to-do lists—is scoped to the extent of their specialized data sets.

The trade-off is pretty clear: A scoped search can be easier to use because the search terms need not be as precise. That is, the user does not have to find *exactly* the right search terms that would unambiguously select out the target from the entire information store. The user can use terms that would be fairly common (and therefore low precision) in an unscoped search, but that would be uncommon (and precise) in a scoped search.

This is particularly true for personal information stores, where the user's metaknowledge of the information store can allow for very effective use, because the user knows about the relative frequency of terms. For instance, a search such as [Diane] that is scoped (limited) to just a personal store of e-mails will pull up all of my e-mails with "Diane." As a general query, [Diane] is too broad a question, but the user happens to know that there only two

Diane's who ever send mail, so it is in fact a reasonable approach to finding a specific e-mail.

By contrast, searching over a fairly heterogeneous information store (e.g., the personal desktop) *broadens* the search to encompass data objects over many different stores.

Traditionally, many personal information applications have created their own private, uninspectable data storage forms. Calendars and contact managers have usually maintained information in their own databases for performance reasons. As a side effect, indexing these databases has generally been difficult from other applications, requiring special purpose format convertors to transcode the internal data into accessible content.

More recently, though, the trend has been toward increasing openness in data storage formats. The consequence of this has been to make more content available to general purpose search engines that can crawl and index a large personal workspace. This lends itself to searches that cut across application boundaries.

Broadly operating search tools such as GDS, X1 desktop search or Apple's Spotlight let the user pose queries that apply across multiple scopes, crossing many application types and storage formats.

There is a clear trade-off: A specialized tool can more easily support complex queries that rely on a detailed knowledge of metadata and content structure, but broadly operating search tools provide exactly the breadth and coverage lacking in specialized applications.

These two approaches to the same problem point out that no single solution seems right for all personal information search tasks. Each approach has its benefits and costs: In many personal information management cases, it seems as though multiple views of data approach might be best. For example, X1 display (Fig. 2) shows a high density of results with preview, while the Google display has text snippets with a thumbnail views by document type (Fig. 3), the Apple Spotlight display (Fig. 4) lists document title and date, with buttons to allow for additional metadata filtering and sorting.

Such broad desktop search tools index e-mail and a wide variety of file types. X1 also provides tabs for entries in the desktop contact list, and IM transcripts, while Google offers web history search. It is clear that desktop search has become much more than just providing a full-text index of ordinary documents—it is a general mechanism for finding content, often personal content, that can be scattered over a variety of source, places and filing systems.

Apple also has a desktop search mechanism—Spotlight—that has a large number of sort-by and filtering capabilities. It too provides capabilities for scoping documents by file type and sorting to get to the right information quickly.

As these three desktop search tools show, there is a wide variety in the way found objects can be displayed to the

Fig. 2 The X1 desktop search application is invoked as a background application, typically via a hotkey. It indexes a wide variety of file types, providing previews of the documents (in the right-hand panel, here showing the Google Desktop icon) and letting the user sort and refine the query
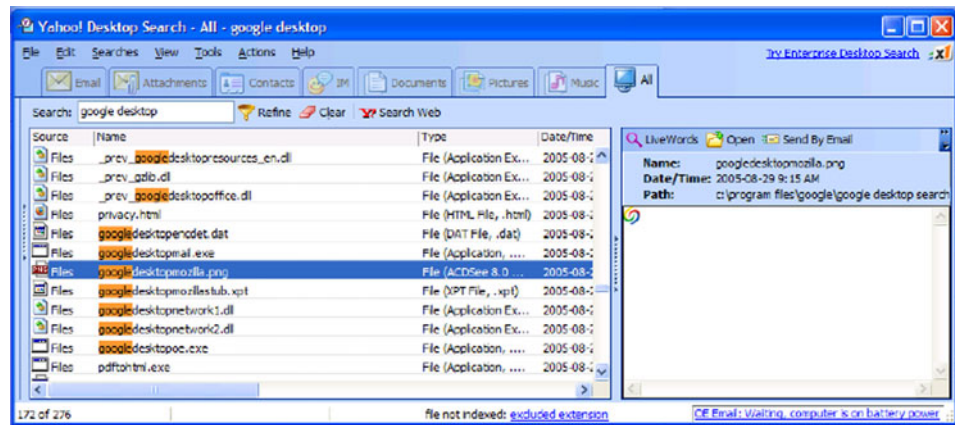

Fig. 3 Google desktop search can also be invoked from an ordinary desktop icon. When started up this way, GDS is limited to searching just the desktop and does not reach out to the web as a whole. Note that e-mails, file, web history and chat transcripts are all searched. The user interface of "just desktop" search differs from that shown in Fig. 1 to reflect this differently scoped search. The Google logo that includes "Desktop" and "Desktop" is incorporated just below the blue separator line. Scoped search (just e-mail, just files, etc.) can be selected by clicking on one of the links in the Desktop: line
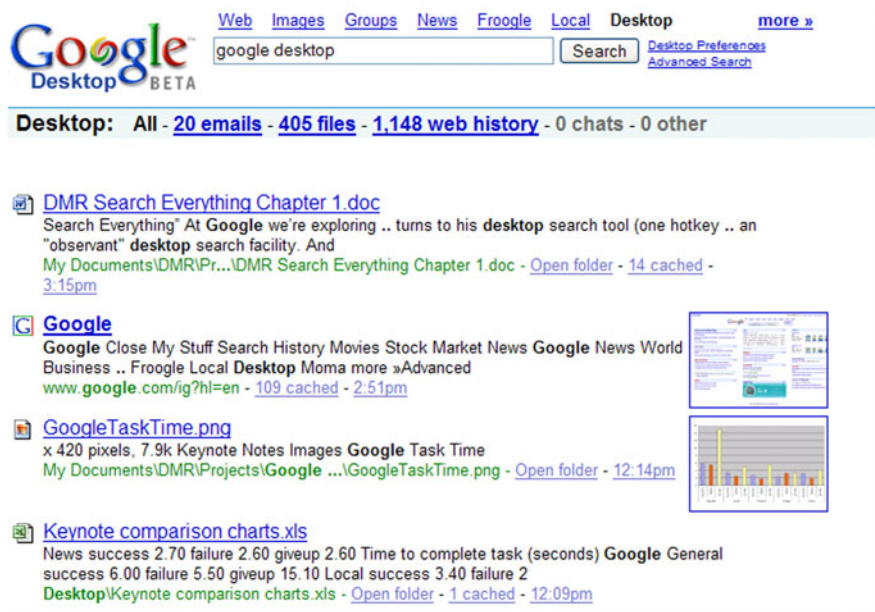
searcher. X1 and Apple desktop search provide different sort orders by object metadata (date, filetype, etc.), where Google has different file type views (all, e-mail, files, web history). Google desktop search provides a short snippet summarizing the document with a thumbnail, while X1 desktop search provides a document preview in the right-hand side pane.

3.3 Access everywhere: toolbars and the future of integration

Just as there has been a trend toward an opening up of content storage formats, allowing the growth of broad search tools, it is believed that there will also be a continuing trend toward allowing search to become a commonplace feature of all applications.

Not only will search be a basic capability (as common as, say, text editing in dialog boxes), but search will be both scoped and broadened simultaneously as it takes place across a wide variety of places within the desktop environment. It is expected that search toolbars and desktop widgets (small applications embedded in the desktop) will proliferate, providing search access for specialized tasks (such as global search from within an application), and universally by providing rapid access to search from anywhere. Information that was previously difficult to find (because it was in a special data repository with awkward, unmemorable access methods) suddenly becomes easily

Fig. 4 Apple's Spotlight desktop search tool is easily invoked by clicking on an ever-present button on the desktop. The right-hand side panel allows grouping of search results and sorting within the group on different properties of the documents. Here, personal information about kid's spring camps is intermixed with bookmarks, tax returns and news
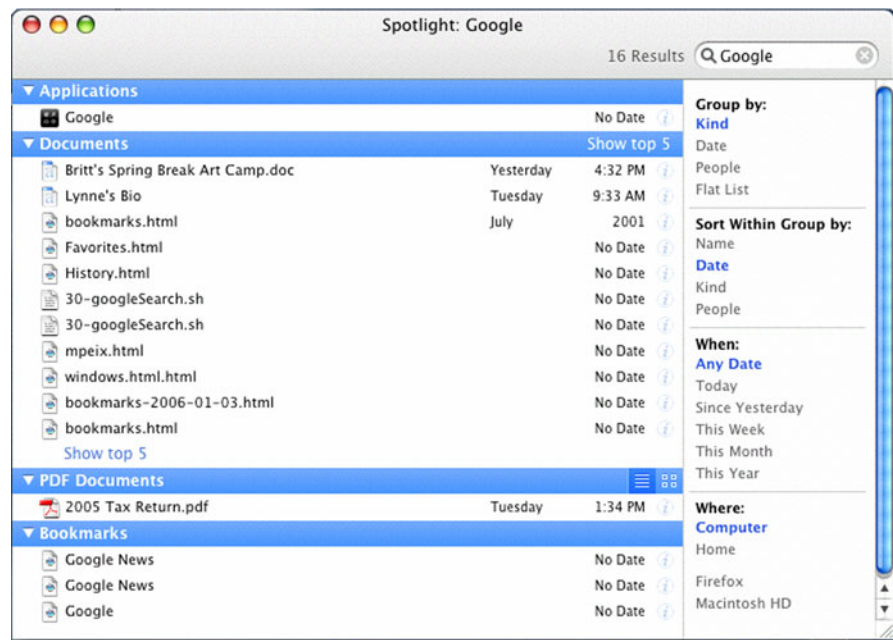
available with a specialized desktop search widget that provides a single-purpose, scoped search.

And just as search continues to expand to provide key access methods for personal information, personal information is becoming increasingly scattered across multiple devices. Many knowledge workers have multiple desktops, and the ability to think of a personal "information cloud" as the aggregate of all personal information of a user (regardless of what devices currently carry it). It is clear that the ability to merge search results over multiple indexed desktops provides a clear path forward for a universally distributed search mechanism over all of one's remotely distributed working content; an effort that is clearly in line with the larger vision of the distributed, collaborative and smart information spaces of the future [23].

nouns, dates and place names seems feasible in the not-too-distant future. Such capabilities will enable smooth integration between different kinds of information streams within one's personal data space. Calendars and calendar entries are becoming indexable and useful across devices and data sets [6].

Just as personal search tools allow integration across multiple personal computers, including other kinds of personal devices seems to be equally critical. Of all the information devices in your personal space, should all be searchable from any device? Coordination between multiple devices is still an open problem. But the future path is clear—a single user "data cloud" will be accessible and searchable from any personal device, with synchronization happening automatically in the background.

4 Looking forward

Innovation will continue to appear in the UI arena—but also in the opportunity for mashups (that is, unexpected combinations of data resources and visual presentations, often web-based) of different content types and different kinds of contextual information.

Advances will continue to be made in text- and web-document-indexing, with newer analytic capabilities providing new kinds of indices and retrieval capabilities. Currently, there is relatively little higher-order recognition of objects in documents. While full-text *understanding* is still a ways off, the ability to recognize and identify proper

5 Conclusions

The goal of personal information management is to let users search all of their information assets—desktop content, personal notes, web, information streams, to-do lists, calendar information—the list continues to grow as new information types become commonplace.

Yet, the growth of new media types is a trend that shows no signs of slowing down. From the vantage point of early 2011, an impressive number of new media types have appeared over the past 18 months: video blogs, photo blogs, instant message transcripts, tags on files, mashed-up maps and data of many different kinds. The list is

constantly evolving with increasing bandwidth, new media mashups and editing/composition systems. Future users will want to be able to search all the content they identify as personal, over data types and media content that are not yet anticipated.

The shift from structured query tools to more free-text forms was not without a loss. Performance is still an issue when crawling and indexing large, semi-structured data for broad use. Users will continue to want increasingly efficient access to a larger number of content types, with some deeper kinds of computed analysis over the content. Increasingly, the ability to refer to dates, people and places in subtle and sophisticated ways will become part of expectations of information management.

And in the limit, the pursuit of ever more personal information will plausibly extend beyond computational desktops and obvious data devices (phones, PDAs, etc.) and into physical objects tagged with RFID and spatial location systems. The extension of personal information management from phone numbers to tracking, locating and searching for objects in personal space does not seem all that far away [3, 21, 25].

Clearly, another implication is that as the virtual workplace grows increasingly distributed across many devices, an understandable model of how personal content is available via cloud-based services will be required. As the underlying information data storage model grows ever less relevant to search tasks, the larger problem of portraying what is within the user's personal information store will become a key feature in ubiquitous information search systems design.

References

1. Anderson, J.: Cognitive Psychology and its Implications. WH Freeman, San Francisco (1991)
2. Apple Spotlight (2006) Desktop Search Application, http://www.apple.com/macosx/what-is-macosx/spotlight.html [last checked: April 15, 2011]
3. Becker C, Dürr F (2005) On location models for ubiquitous computing. Personal and Ubiquitous Computing, 9(1): 20–31
4. Bellotti, V., Ducheneaut, N., Howard, M.A., Smith, I.E.: Taking email to task: the design and evaluation of a task management centered email tool. In: ACM Conference on Human Factors in Computing Systems Fort Lauderdale, pp. 345–352. ACM, FL (2003)
5. Bush, V.: As we may think. Atl. Mon. **176**(1), 101–108 (1945)
6. Collins, M.: Ranking algorithms for named-entity extraction: boosting and the voted perceptron. In: Proceedings of the 40th Annual Meeting on Association for Computational Linguistics, pp 489–496. Philadelphia, PA (2001)
7. Cutrell, E., Dumais, S.T., Teevan, J.: Searching to eliminate personal information management. Commun ACM **49**(1), 58–64 (2006)
8. Czerwinski, M., Gage, D., Gemmel, J., Marshall, C.C., Perez-Quinones, M., Skeels, M.M. et al.: (2006) Digital memories in an era of ubiquitous computing and abundant storage. Communications of the ACM, Special Issue on Personal Information Management **49**(1):44–50
9. Davies, G., Thomson, D. (eds.): Memory in Context: Context in Memory. John Wiley & Sons, Inc., Chichester (1988)
10. dtSearch, Desktop Search Application http://www.dtsearch.com [last checked: April 15, 2006]
11. Ducheneaut, N., Bellotti, V.: Email as Habitat: An exploration of Embedded Personal Information Management. ACM Interactions. 2001 September-October; 30–38 (2001)
12. Dumais, S., Cutrell, E., Cadiz, JJ, Jancke, G., Sarin, R., Robbins, D.C. (2003) Stuff I've seen: a system for personal information retrieval and re-use, Proceedings of the 26th annual international ACM SIGIR conference on Research and development in information retrieval, Toronto, Canada
13. Google, Google Desktop Search (2006) http://desktop.google.com/ [last checked: March 16, 2011]
14. Jones, W., Phuwanartnurak, A.J., Gill, R., Bruce, H.: Don't take my folders away!: organizing personal information to get things done, CHI '05 extended abstracts on Human factors in computing systems. Portland, OR (2005)
15. Kidd, A.: The marks are on the knowledge worker. In Proceedings of CHI'94 Human Factors in Computing Systems, pp. 186–191. ACM Press, New York (1994)
16. Lansdale, M., Edmonds, E.: Using memory for events in the design of personal filing systems. Intl. J. Man-Mach. Stud. **36**(1), 97–126 (1992)
17. Lansdale, M.: The psychology of personal information management. Appl Ergon **19**(1), 55–66 (1988). v 3, 25
18. Malone, T.: How do people organize their desktops? Implications for the design of office information systems. ACM Trans. Office Inform. Syst. **1**, 99–112 (1983)
19. Morville, P.: Ambient Findability. O'Reilly Press, Sebastopol (2006)
20. Nutch (2006) http://lucene.apache.org/nutch/ [last checked: May 15, 2011]
21. Raper, J., Gartner, G., Karimi, H., Rizos, C. (2007) A critical evaluation of location based services and their potential. J. Locat.-Based Serv **1**(1)
22. Streitz, N., Magerkurth. C., Prante, T., Röcker, C.: From information design to experience design: smart artefacts and the disappearing computer. ACM Interactions, Special Issue on Ambient intelligence (July & August 2005), **12**(4):21–25 (2005a)
23. Streitz, N., Nixon, P.: The disappearing computer. Commun. ACM **48**(3), 33–35 (2005)
24. Streitz, N.A., Kameas, A., Mavrommati, I. (eds.): The disappearing computer. Lecture Notes in Computer Science, vol. 4500. Springer, Berlin (2007)
25. Want, R., Russell, D.M.: "Tags, IDs and what to do with them" Distributed Systems Online, **1**(1) (Sept 2000)
26. Whittaker, S., Sidner, C.: Email overload: exploring personal information management of email. In: CHI '96: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 276–283. ACM Press, New York (1996)
27. Whittaker, S., Bellotti, V., Gwizdka, J.: Email in personal information management. Commun. ACM **49**(1), 68–73 (2006)
28. X1, X1 Desktop Search (2011), http://www.x1.com/products/ [last checked: May 15, 2011]