

# 10. Everything through Email

STEVE WHITTAKER<sup>1</sup>, VICTORIA BELLOTTI<sup>2</sup> AND JACEK GWIZDKA<sup>3</sup>

## 10.1. Introduction and Description of Topic

Email is one of the most successful and frequently used computer applications. According to marketing surveys, over 4 billion corporate email messages were exchanged per day in 2001, increasing to a projected 35 billion in 2005. Moreover, 97% of workers report using email multiple times per week for a daily average of 49 minutes (Gartner 2001, Levitt, 2000, Pitney Bowes, 2000). For many people, work is interpersonal rather than solitary, and email is the main conduit through which their work and information are distributed. They tend to 'live' in email, as evidenced by the sheer amount of time they spend using it (Bellotti, et al., 2005), and their evaluation of its importance, with 71% of people stating that it is "essential" for their everyday work (Pitney Bowes, 2000).

Email serves as an information conduit – acting as a delivery channel for different types of information, including documents, slides, contact information, and schedules. Its use as conduit naturally leads it to be used for key PIM functions. People use their inboxes as to-do lists to manage current tasks, their email folders as a repository for archival information, and their email address books to find contacts. They even use email to find and schedule calendar appointments.

Despite its success, there are significant problems with email. Users complain about feeling overwhelmed by the volume of messages they receive, and in processing incoming messages effectively (Bälter and Sidner, 2002, Bellotti et al., 2005, Venolia, 2001, Whittaker and Sidner, 1996). They have difficulties, too, in organizing and managing archives (Bälter and Sidner, 2002, Boardman and Sasse, 2004, Whittaker and Sidner, 1996). They experience severe problems using email to manage tasks, leading them to forget tasks and obligations (Bellotti et al., 2003, 2005, Whittaker et al., 2002a, Whittaker, 2005). Email's failure to address these PIM problems threatens to seriously reduce individual and corporate productivity.

Why is email so hard to process and manage and why does it generate such serious PIM problems? And yet why is it still so popular and so essential in PIM? And what can be done to

---

<sup>1</sup> University of Sheffield, s.whittaker@shef.ac.uk.

<sup>2</sup> Palo Alto Research Center, bellotti@parc.com

<sup>3</sup> Rutgers University, pim.bch@gwizdka.com

make it better able to support this essential role? This chapter will illustrate how many of email's PIM problems stem from its conduit function. Email is incessant, because it is often the primary delivery channel for work and information (Duchenaut and Bellotti, 2001, Venolia et al., 2001, Whittaker, 2005). It is hard to process and organize because it is a mixture of different types of information (tasks, documents, FYIs, meeting scheduling), some of which are highly relevant (work tasks) and others irrelevant (jokes). And most email is generated by others - making it harder to understand, evaluate and organize than personally generated information (Boardman and Sasse, 2004). These problems are exacerbated by the fact that most email systems have no inbuilt support for PIM besides folders, so that users have to devise ad hoc ways to manage tasks, find contacts and organize useful information (Bellotti et al., 2003, 2005, Whittaker, 2005, Whittaker and Sidner, 1996).

This chapter discusses how people use email to manage their personal information, and describes the tools that can support such behavior. We look at important aspects of PIM through the lens of email: allocating attention, deciding actions, managing tasks, and organizing messages into folders.

## 10.2. Email Activities and Their Relation to Finding, Management and Keeping Aspects of PIM

Figure 1 shows major email activities and the relations between them. We now describe these activities and identify the PIM problems that each creates. In each case we discuss the extent to which these problems overlap with problems occurring in other areas of PIM.

1. *Allocating Attention.* Email messages differ in their importance, urgency and interest (Dabbish et al., 2005). They also take differing amounts of time and effort to process. As a result, users do not access email messages sequentially, instead deciding which messages they attend to first (Bälter and Sidner, 2002, Bellotti et al., 2005, Venolia et al., 2001).
2. *Deciding Actions.* Unlike other information items obtained from the web or document repositories, email messages may communicate expectations from others of some kind of response. Users have to assess whether this is the case and decide what action to take for a given message. This is similar to, but more complex than the classic PIM keeping decision (Bellotti et al., 2005, Dabbish et al., 2005, Venolia et al., 2001).
3. *Managing Tasks.* Many email tasks are not one-shot and cannot be discharged in a single session (Bellotti et al., 2005, Gwizdka and Chignell, 2004; Whittaker, 2005). Task management is complicated because users have to remember to execute tasks, which often require complex collation of information from multiple sources and other people. Task management relates to refinding activities. Users need to structure personal information to guarantee it can be found easily and at the appropriate time.

4. *Organizing Messages and Folders.* As in other aspects of PIM, users have to decide how to organize information. This relates to keeping and management activities. If users decide to file, which folder should they put it in? Again these decisions are complicated by the fact that much email information is generated by others, sometimes without the context of personal user goals or initiatives. This in turn means that folder organization is often unsuccessful (Bälter 1998, Whittaker and Sidner, 1996).

Note that the above list does not imply a strict temporal sequence, with messages first being attended to, decisions made, tasks managed and then filed. As shown in Figure 1, email processing can take different forms depending on the nature of the message, its relation to other tasks, and user workload. Simple irrelevant messages may be scanned and deleted immediately. Other messages may be left in the inbox, returned to several times as the user scans their inbox, later judged to be irrelevant and deleted. Yet other messages may be scanned and filed if they relate to ongoing projects or interests but do not demand a response. Messages related to interdependent tasks, may be filed or left in the inbox as reminders, they may be repeatedly revisited as the task unfolds, or the user encounters them in the context of accessing their inbox. And older archived messages may be re-accessed if they have relevance for current tasks.

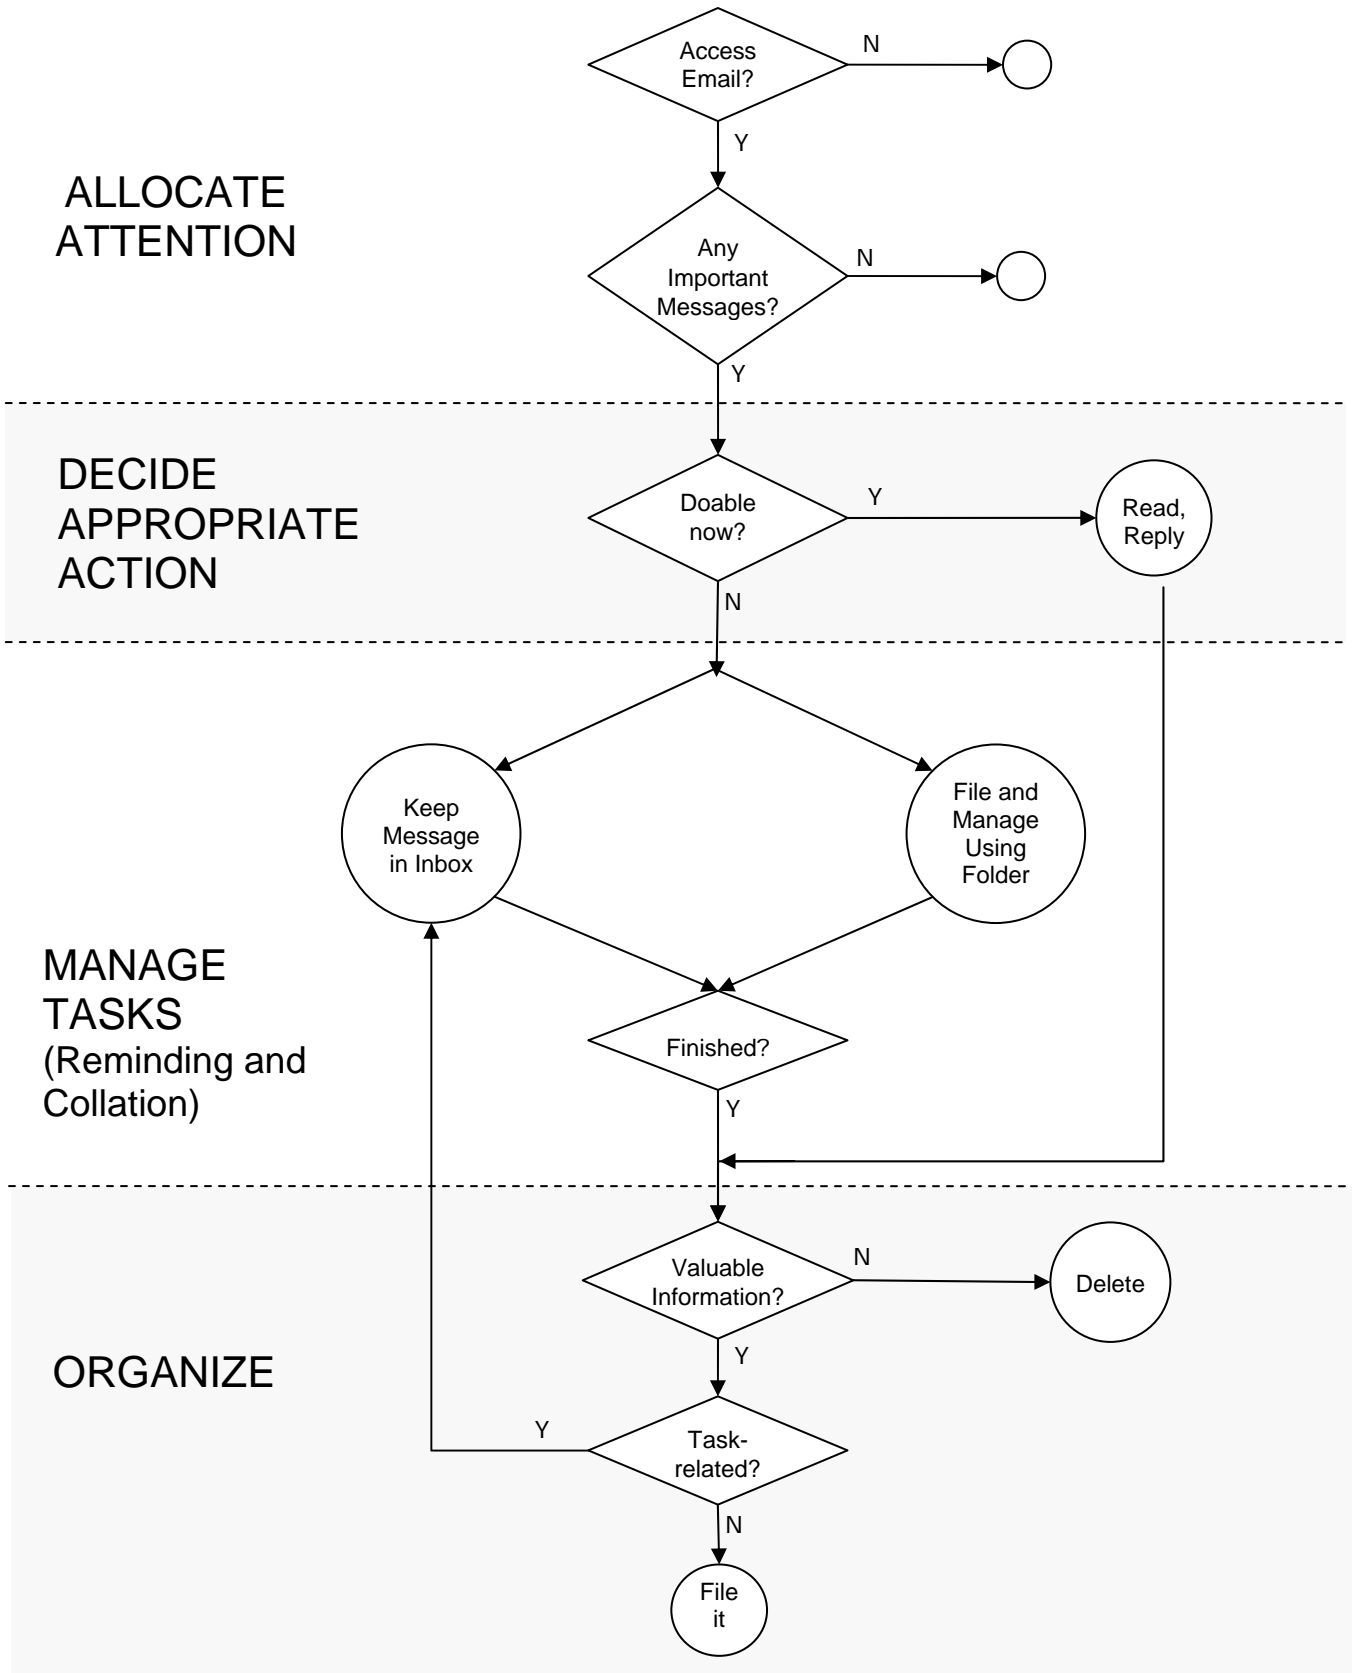


Figure 1. Email processing.

## 10.3. Understanding Email Tasks

We now review research of special relevance to email's role in PIM. To understand the challenges of managing email, it is important first to document its general properties. We present some general statistics on email volume, inbox and folder characteristics, followed by a taxonomy of the different types of email message, and how people spend their time processing email.

### 10.3.1. Scoping the Problem: Some Basic Email Properties

Table 1 summarizes the findings from 8 studies of email averaged over all the subjects in all the studies (between 86 and 244 individuals for each figure).

Data	Number of Msgs Received/Day	Number of Msgs Sent/Day	Number of Inbox Msgs	Percentage of Total Msgs Kept in Inbox	Total Number of Email Msgs Retained	Number of Folders
#Participants	86	200	206	183	183	244
Weighted Average	43.7	13.7	1424	49.3	2846	38.8

Table 1: Summarized Basic Email Statistics. Sources: Mackay (1988), Whittaker and Sidner (1996), Bälter and Sidner (2002), Ducheneaut and Bellotti (2001), Venolia et al (2001), Boardman and Sasse (2004), Bellotti et al. (2005), Dabbish et al. (2005).

Table 1 summarizes 8 studies that have collected general data about email behaviors. One important proviso in interpreting the data is that each of the reported studies reported large individual differences between users, depending on their email processing strategies - so there are large standard deviations for the data presented. Individual differences in PIM are discussed in detail in Chapter XX. Note also that these studies have focused on workplace rather than recreational usage and the specific context of the workplace probably influences both the volume of emails and the manner in which they are processed.

Somewhat surprisingly, Table 1 does not suggest a huge increase in email loads over the last 18 years. This may be because these studies researched populations working in high technology or academia where email throughput has always been high, and is perhaps closer to a natural saturation point.

Nevertheless, certain general observations may explain users' complaints about processing and managing email. For example, people have to process a large number (between 38 and 52) of messages each day. They send 31% as many emails as they receive – suggesting many messages they receive do not demand a response. On the management side, they keep hundreds or even thousands of messages in their inbox, accounting for about half the total messages they keep. With the proviso that there are large individual differences in how people organize information,

these statistics can be explained by the use of the inbox as a ‘to-do’ list for working tasks. On the organizational side, they create on average 38.8 folders.

### *Message Types*

There have been various attempts to categorize different types of messages that people receive. For example Whittaker and Sidner (1996) describe messages as ‘to-dos’, ‘to-reads’, ‘messages of indeterminate status’ and ‘ongoing correspondence’, based on properties such as message age, length, and whether or not they are filed or left in the inbox. They also present some prevalence statistics, with ‘to-reads’ accounting for 21% of the inbox. Dabbish et al., (2005) present a more systematic analysis of message types in a study of 144 users who were asked to select 5 recent email messages they had received and to classify these into various speech act types. Messages could serve multiple functions, but the overall breakdown was Action Requests (34%), Information Requests (18%), Information Attachments (36%), Status Updates for Projects/Tasks (21%), Scheduling Requests/Responses to Scheduling Requests (14%), Reminders for Meetings/Actions (16%), Social Messages (8%), Other (12%).

An important point to make about this data is that many messages require the user *to act in some way* – whether this is to carry out an action, respond to a request or to schedule a meeting. Such *interdependence* has implications for how effectively others carry out their tasks. Section 1.3.4 details the problems email presents in managing tasks.

Finally, there have been video studies of how people allocate time processing email (Bellotti et al 2005):

- 54% composing messages
- 23% reading messages, attachments and links
- 10% filing messages
- 6% scanning inbox messages for things to read or deal with
- 2 % deleting messages
- 2% looking for messages in folders
- 2% managing attachments

In addition to the obvious dominant activities of reading and writing messages, users dedicate proportionally large amounts of time filing messages. In Section 10.3.5 we examine why filing is such a demanding activity and evaluate its success in creating useful structures for future retrieval. Section 10.3.2 and 10.3.3 discuss how new emails are processed and Section 10.3.4 explains how the inbox is used for task management.

## **10.3.2. Allocating Attention**

There are two parts to attention allocation. The first involves *when and why* people decide to attend to email, and the second *how* they decide which messages to attend to, and in what order. Allocation of attention is influenced by multiple factors including user workload, available time, nature of the message, identity of its sender, and so on. This task is somewhat related to the typical PIM task of finding information (Barreau and Nardi, 1995; Lansdale, 1988). Finding is typically a two-phase generate and test process where users first identify a set of potentially

relevant information items, either by browsing or search. They then decide which of that set is relevant to their current activity. Allocating attention in email is similar to the second phase of the process.

*Deciding when to attend to email.* Despite email's importance for task delegation, there has been relatively little systematic study of what causes people to attend to their email. Mackay (1988), and Whittaker and Sidner (1996) note how email dominates people's work lives. Similar findings led Duchenaut and Bellotti (2001) to make the observation that people 'live' in email – leading them to characterize it as a 'habitat'. In Mackay's (1988) study, 13/27 informants stated that they read email 'constantly'. However as Venolia et al (2001) observe, people cannot spend their entire work time in email, raising questions about how often they access email, as well as the factors that lead them to switch to email from other activities. Dabbish et al. (2005) report that people check email 19 times/day, so what prompts this? For some, any incoming unread message serves as a trigger. Venolia et al. present survey data from 406 Microsoft employees suggesting that most of their respondents attend to incoming messages as soon as these arrive. Participants also usually access email as soon as they can after they have been away from it, e.g. first thing in the morning, or upon returning to their desk after an absence. And because email serves as a task manager, archive and contact manager, email access may be prompted by the need to retrieve 'to-dos', reminders or archived information relevant to a current task (Duchenaut and Bellotti (2001), Whittaker and Sidner, 1996). Access frequency is directly influenced by task management demands; Gwizdka (2004b) reports that the more task-related messages users keep in the inbox the more likely they are to constantly read email, rather than at specific times only.

*Deciding which messages to attend to.* Users generally don't access messages sequentially in the order in which they arrive, and they often don't access all messages. Bälter and Sidner (2002) observed that their 10 users scanned the inbox a mean of 2.3 times to decide which messages to focus on, and Venolia et al. report that 70% of their interviewees processed messages out of sequence. Bälter and Sidner also found that only half their informants accessed all the emails they received. But how do people decide which messages to access first and which to ignore? One possibility is that people focus on *important messages*, and Venolia et al. (2001) list a variety of factors that affect perceived importance, such as the identity of the sender, number of recipients, the nature of the message and whether they message is a reply. Dabbish et al. found that important messages were those that demanded action, conveyed task status, discussed schedules or were reminders. Important messages were also more likely to be sent by people with whom the recipient had a strong working relationship, and to be personally addressed rather than to a wide distribution list. They were also highly unlikely to be social messages.

### **10.3.3. Deciding Actions**

Having decided *which* messages to attend to, users then have to decide *what* to do about them, for example whether to respond, file, delete, or defer action until later. In other areas of PIM, users have to make decisions about the information that they encounter, such as whether to keep it or not (Jones, 2004). In email these decisions are more complex, and they may also have important implications for the work of others, e.g. when the information concerns an interdependent task.

Dabbish et al (2005) found that people only respond (or intend to respond, in cases where they defer action) to about 35% of messages that they receive. Contrary to their expectations, Dabbish et al. also discovered that perceived importance is only weakly correlated with responding; people are only 8% more likely to respond to important messages. In contrast, sender characteristics (15%), whether the message is personally addressed (20%) and, most strikingly, whether the message is social (23%) are all stronger determinants than importance in predicting responding behavior. Another notable strategy is that users defer responses to 37% of messages they intend to respond to. There are several possible reasons for this including, time for reflection, or with interdependent tasks, communicating with others to gather and collate the information needed for a collective response. Whether people respond to a message immediately or not, their most likely action for a message requiring a response is to leave it in the inbox. Even when they have already replied to the message, users leave it in the inbox 55% of the time, whereas if the message has not been responded to, they leave it in the inbox 79% of the time. This suggests that people are using the inbox as a reminder to respond – using the inbox for task management as described in the next section.

Dabbish et al. also examined whether incoming messages were filed, deleted or retained in the inbox. Overall people filed 27% of their messages, deleted 24% and kept the remaining 49% in the inbox. People were more likely to keep important messages in the inbox, although message type and sender characteristics did not directly affect retention behaviors. Individual differences were important here – accounting for 48% of the variance in the probability of a message being left in the inbox.

Of the messages that don't demand a reply, users file 28% of (presumably useful) messages, delete a further 30% of (presumably irrelevant) messages, but then retain the remaining 42% in the inbox. Why do they retain so many of these informational messages that don't demand a reply? Some messages may be kept for later reading. Whittaker and Sidner (1996) found that 21% of inbox messages contained more than 5 screenfulls of text. Immediate time pressure may lead people to defer reading these until they have more time. These 'to reads' may not have been filed because leaving them in the inbox is a way to remind users that they still have to be processed. Other messages are retained because users don't know what to do with them. Rather than investing valuable time in reading a new message at once, users register its arrival, but defer dealing with it until they are more certain of its importance. Rather than delete it immediately, they conservatively retain it, in case it turns out to be important.

Messages are often revisited multiple times, with different decisions taken about how they should be processed. For example a message might be kept in the inbox or stored in a folder and accessed multiple times until a task is complete, when it might be filed in an archive (Bellotti et al., 2005).

One obvious conclusion is that deciding what to do with email messages is rather more complex than the typical PIM keeping decision. While users do end up keeping the majority of messages, they make numerous other decisions than a simple keep/delete, e.g. what action to carry out on a message, and whether and when to reply to it. There also seem to be two important and distinct types of keeping activity; (a) filing – which is the typical PIM task of categorizing information



for future use, and (b) inbox retention – where information is not organized, but kept in a working buffer to remind people it may require some future action. The next section describes this complex management activity.

#### 10.3.4. Managing Tasks

A classic PIM activity is that of managing information to meet the demands of specific tasks, and many PIM studies have described how users categorize information in order to facilitate its future retrieval (Jones, Kidd, 1994, Lansdale, 1988, Teevan et al.). Task management in emails turns out to have rather different demands however.

*Task management* involves reminding oneself about current tasks, tracking task status, as well as collating and maintaining information relevant to those tasks. A direct consequence of email functioning as a conduit for work and information is that people use it as a *de facto* task manager. Many email tasks are too complex or lengthy to be executed in one-shot (Bellotti et al., 2005, Venolia et al., 2001, Whittaker and Sidner, 1996, Whittaker, 2005), leading to deferral for 37% of messages that require replies (Dabbish et al., 2005).

Deferral is often a direct consequence of *interdependent* tasks, i.e. those requiring collective activity (Bellotti et al., 2005). Two problematic aspects of managing interdependent tasks result from *iteration* and *delays* between messages relating to the task. Iteration arises because interdependent tasks often require multiple exchanges between participants before they can be resolved (Bellotti et al., 2005, Venolia et al., 2001, 2003, Whittaker and Sidner, 1996). People may need to negotiate exactly what a task involves, or multiple responses need to be synthesized. Delays occur because collaborators lack the necessary information to respond immediately to address their part of the task. One way to estimate the prevalence of interdependent tasks is by determining how many emails are part of a *conversational thread*, as threads indicate relations and common underlying activities among messages. Threading estimates range from 30-62% of messages (Bellotti et al., 2003, Fisher and Moody, 2001, Klimt and Yang, 2004, Venolia et al., 2001)<sup>4</sup>. The largest available public email corpus from Enron has 62% of messages appearing in threads, with a mean thread length of 4.1 messages, though the median length is only 2 (Klimt and Yang, 2004).

Delays and iteration give rise to two critical activities in managing interdependent tasks. Users have to leave themselves *reminders* about outstanding tasks, and to *collate* different elements of the same collaborative task. We now look at how people address these problems, to carry out core task management activities.

---

<sup>4</sup> These threading figures generally underestimate interdependent tasks. They ignore new interdependent tasks delegated to users who have yet to deal with them, and they do not register tasks forwarded to others. Furthermore, people are often casual replying: sometimes failing to reply using *re*. They may occasionally overestimate; e.g. when someone initiates a new task by replying to an old message from another user (Erickson and Kellogg, 2001, Herring, 1999).

People predominantly use email itself to manage tasks arriving in email, rather than relying on a dedicated task management application, such as a 'to-do' list on a PDA (Bellotti et al., 2004, Venolia et al., 2001). New messages relating to a particular task arrive in email, which, as email is constantly processed, serve to remind users about that task, whereas copying information into a separate application requires additional effort and bookkeeping.

One *reminding* strategy involves setting up a dedicated 'to-do' folder containing reminders about outstanding tasks (Bellotti et al., 2003, 2005, Whittaker and Sidner, 1996, Whittaker et al., 2002a). However Whittaker and Sidner (1996) report that 95% of users abandoned this strategy, because it requires an additional cognitive step. Instead of being reminded about tasks as they access new email in their inbox, people have to explicitly remember to open the 'to-do' folder, and they often forgot to do this. Another strategy is to organize emails relating to current tasks into active folders, returning to these folders when they need to deal with those tasks (Bellotti et al., 2005; Gwizdka, 2004b). This has the advantage of grouping messages allowing them to be more efficiently worked on together. But this strategy only works if users develop the habit of routinely returning to inspect those folders, as most users do with the inbox.

But the most common strategy for managing interdependent tasks is to respond or forward the original message to relevant others, leaving the original message in the inbox as a reminder about the task (Dabbish et al., 2005, Mackay, 1988, Venolia et al., 2001, Whittaker and Sidner, 1996, Whittaker, 2005). This serves to manipulate attention. Users know that they will return to the inbox to access new messages. In the course of doing so, they hope that they will see the reminder and recall the outstanding task. Users may even send themselves emails as reminders or containing important information (Jones et al., 2002). This inbox reminding strategy is far more common than other techniques such as using message flags, classifying messages as 'to-do' items, relying on external applications, or external reminders (Venolia et al., 2001; Gwizdka, 2004b).

A second problem with interdependent tasks is *collating* information, i.e. assimilating the disparate pieces of information relating to a common task. Collation is difficult because messages relating to the same activity may come from multiple individuals using different subject lines and even end up in different folders. As a result, some users feel that keeping messages in their inbox makes them more *collatable*, in part because the multiple passes carried out when processing the inbox, serve to remind them about its contents (Whittaker et al., 2005).

However, many users experience severe problems with using the inbox for task management. Reminding and collating are highly dependent on *visual scanning* which becomes less effective as the number of inbox items increases. Most email interfaces present a header list along with message body information for a selected message. Although users differ in how they configure their header list, they tend to operate with 10-20 message headers visible on the screen (Whittaker and Sidner, 1996, Whittaker et al., 2002a). This means, on average, that less than 2% of inbox messages are visible.

Because most inbox message headers are not constantly in view, users have to actively remember to scroll back to view older messages. And it is hard to identify specific, related messages scattered at intervals through the inbox among hundreds of other unrelated items. The

alternatives of searching and sorting by metadata such as sender or subject are only effective when users can remember salient information about the message. For this reason, older inbox messages are often characterized as being “out of sight and out of mind” (Whittaker and Sidner, 1996).

And relying on folders or messages threads is little better: Venolia et al. found that 23% of message threads occurred in more than one folder. Using threads is also unreliable because they rely on a common message subject line – and this turns out to be an inconsistent indicator that messages relate to the same task (Ducheneaut and Bellotti, 2003).

A related problem is *obligation management*, i.e. keeping track of the tasks that you ‘owe’ others or are ‘owed’ to you. Although the strategy of leaving messages in the inbox works reasonably well for reminding about tasks that one receives, it is much less effective when delegating tasks to others. Most email programs keep a copy of sent messages, but these are generally stored in a separate ‘sent’ folder. Having to remember to access the ‘sent’ folder suffers from the same weaknesses as using a ‘to-do’ folder. Some users attempt to avoid this by copying themselves on every sent message to generate inbox reminders, but this overloads the inbox with duplicate messages for every task (Bellotti et al., 2005, Whittaker et al., 2002a, Whittaker 2005).

In conclusion, task management is problematic in email. Although the default strategy is to use the inbox as a ‘to-do’ list, there are numerous problems associated with doing this, including overlooking reminder messages, and inability to collate information. These problems do not have clear analogues elsewhere in PIM.

### **10.3.5. Organizing Messages into Folders**

People dedicate a great deal of time and effort to organizing their email. Many PIM studies have documented the problems of categorizing information in a way that facilitates later retrieval (Barreau and Nardi, 1995, Boardman and Sasse, 2004, Jones et al., 2002, Lansdale, 1988, Whittaker and Hirschberg, 2001). Email presents similar organizational problems.

Creating email folders is a common activity, with users overall filing 27% of their messages, although there are large individual differences in the complexity of folder organization. Some users create organizations that include hundreds of folders with complex internal structure; others create few folders (Boardman and Sasse, 2004, Venolia et al., 2001, Whittaker and Sidner, 1996). One important reason for filing is that messages contain useful reference information such as documents, presentations or web links (Bellotti et al., 2005, Ducheneaut and Bellotti, 2001, Venolia et al., 2001). For example, Dabbish et al. (2005) estimate that about 36% of messages contain useful reference information. Indeed some people tend to use email as a document archive, sometimes using metadata such as sender and date to organize and version different document drafts (Bellotti and Smith, 2000). And contrary to Barreau and Nardi’s (1995) assertions, users do access old messages and value their archives (Boardman and Sasse, 2004, Venolia et al., 2001).

Various studies have examined the *types* of folders that email users create. The most common types of folders relate to work projects, contacts and mailing lists (Boardman and Sasse 2004).

There have been fewer studies of the process of creating folders however, although Whittaker and Sidner (1996) examined the problems users experienced with this and we summarise their findings here.

Filing is a cognitively difficult task (Lansdale, 1988). Successful filing is highly dependent on being able to anticipate future retrieval requirements. It is hard to decide which existing folder is appropriate, or, if a new folder is needed, how to give it an appropriate and memorable name.

Users may not file messages because failing to remember where information has been filed could be disastrous. And we have already seen that another reason for not filing is to postpone judgments, in order to determine the value of information. Users want to avoid archiving information that later turns out to be useless or irrelevant.

Even when users do decide to file, folders may not turn out to be especially useful. Users may be unable to remember folder names, especially after a time has elapsed, or when they have large numbers of folders. They have to remember the definition of each when filing and to be careful not to create new folders that duplicate or overlap with existing ones. Duplication detracts from their use in retrieval.

In addition, folders can be *too small* to be useful. A major aim of filing is to collate the huge number of undifferentiated inbox items into a relatively small set of folders containing multiple related messages. Filing is clearly not successful if each folder is small: a folder containing few items has not significantly reduced the complexity of the inbox, nor collated significant amounts of related material. However Whittaker and Sidner (1996) show that filing often fails: 35% of users' folders contain only one or two items. Furthermore, the user has the dual overheads of (a) first creating these folders, and (b) remembering multiple definitions every time there is a new filing decision. Quantitative data illustrate the problems of trying to remember multiple folder definitions. The more folders a user has, the more likely they are to generate "failed folders" containing only one or two items.

Folders can also fail because they are *too big*. When there are too many messages in a folder, it becomes unwieldy. It is difficult to find relevant message in a large folder, as the relationships between different messages in the folder become tenuous, and one benefit of keeping them together is much reduced.

There has been less study of how users access their email archives. Some studies claim that users report few problems in locating information in their archives. However, Boardman and Sasse (2004) found otherwise: users experienced considerable problems when asked to find archived messages. The most common strategies for accessing archival information are to browse folders or to sort messages by sender or subject. Search is much less commonly used, although this may reflect the inefficiency of search in many email applications, and recent developments in desktop search should change this (Boardman and Sasse 2004, Dumais et al., 2003, Venolia et al., 2001). Dumais et al. (2003) looked at common patterns of search over archives, finding that recency was a very strong predictor that a message would be accessed, and that access by sender or recipient name was another common search strategy.

In summary, various studies have shown that in common with other areas of PIM, email users experience problems in organizing their information. They spend large amounts of time constructing structures for future retrieval that may not be greatly effective in supporting access.

## 10.4. Techniques to Support PIM in Email

We now discuss how various technologies address these email activities and PIM problems. From our prior discussion it is clear that there are strong interrelations between different email activities. Thus, a particular technology such as automatic filing may help with both organization of email, and with task management; by reducing inbox size it allows users to better focus on outstanding messages. And techniques that analyze message content and headers might both help users to allocate attention to important messages and to decide what action to apply to a message. Table 2 shows a breakdown of different techniques and the PIM activities they support.

Table 2. Techniques for Supporting Key Email Activities

Technique		Email Activities Addressed
Spam filtering		Attention Allocation
Personal filtering	sender specified	Attention Allocation, Deciding Actions
	user created	Attention Allocation, Deciding Actions
	machine learning	Attention Allocation, Deciding Actions
	pre-defined rules	Attention Allocation, Deciding Actions
AI message summarization		Deciding Actions
Information structuring		Task Management
Embedded task management support		Task Management
Workflow systems		Task Management
Assisted filing		Archive Management, Task Management
Search		Archive Management

### 10.4.1. Allocating Attention

Two main approaches have been taken to help users allocate attention more effectively. Both rely on filtering, but have different aims. *Spam filtering*, aims to remove irrelevant information. *Personal filtering* takes the opposite approach in attempting to identify, highlight or prioritize important information in the inbox.

*Spam filtering*: Spam has become a major problem for most email users, with AOL estimating that it blocks 1.4M spam messages each day – an average of 22 per user per day (AOL, 2003). Many organizations find it necessary to apply filters at the email server to remove suspect content. Spam detection has been relatively successful (if not yet completely accurate) because spam messages often have distinct properties such as large distribution lists, predictable headers

and somewhat predictable message content (Cranor and LaMacchia, 1998). But the race between those that develop spam filters and the spammers evolving techniques to defeat them remains an ongoing contest, with recent spamming techniques involving the generation of more plausible message content.

*Personal filtering:* Personal filtering relies on various techniques depending on how filtering rules are created. One technique relies on the sender to specify information that allows the user to better judge message utility. For example Information Lens (Malone et al., 1986) or the Coordinator (Flores et al., 1988) allow the sender to specify the priority of the message or its projected action. Another technique used by many commercial email clients allows the user to define rules to prioritize incoming messages, using information about the sender, recipients, message content, subject line, etc.. Yet another (predominantly research) approach is to use machine learning to automatically induce prioritization rules. Horvitz et al., (1999) and Metral (1993) built systems that track user triage actions to learn the properties of messages the user attends to, in what order, using this information to identify new messages that are likely to be important. Finally, systems can provide pre-defined rules which users can tune to meet their needs. For example, Marx and Schmandt (1996) use information from calendar appointments, outgoing messages, phone calls and the user's rolodex to infer message importance. Bälter and Sidner (2002) propose a much simpler set of predefined rules that identify urgent messages, important senders, and personal messages.

Personal filtering has not been successful however. It is not commonly used despite being prototyped almost 20 years ago, and being available for several years in commercial products (Bälter and Sidner, 2002, Bellotti et al., 2003, Whittaker and Sidner, 1996, Whittaker et al., 2005). One problem is that writing filtering rules is a programming task - which most users find both hard and time-consuming. And often the users who are most in need of filters have the least time to dedicate to writing rules. Systems that provide predefined general rules simplify the programming task, but this approach has not been successful either. Although machine learning and bespoke filters do not require users to write rules, trust is a more serious concern with filtering approaches. Users are unconfident that rules will operate correctly (Pazzani, 2000). In particular, they worry that mis-defined rules will lead important messages to be overlooked (Whittaker and Sidner, 1996). And whatever the approach to defining filtering rules, these must be modified frequently to reflect changes in people's work activities - making rule maintenance an important issue (Whittaker and Sidner, 1996). In sum, while personal filtering may work for restricted message types with predictable characteristics (e.g. originating from specific newsgroups or users), this technique has limited utility for many messages that users receive.

## **10.4.2. Deciding Appropriate Message Actions**

Recent research has developed techniques that suggest actions for new messages by analyzing past user actions on similar messages. As much of this work is at a preliminary stage, and we postpone analysis to section 1.5.3 where we discuss future email trends.

## **10.4.3. Managing Tasks**

People have proposed three approaches to task management: (a) *information structuring* which seeks to impose organization on the inbox by clustering together messages related to the same task, (b) *embedded support for task management* – where techniques are used to recognize and organize task related information building in direct support for reminding and collation, and (c) *workflow* which seeks to organize messages in terms of underlying task sequences and the roles of the people executing the task.

*Information Structuring.* Most *information structuring* work proposes novel visualizations for inbox message threads, imposing order on the undifferentiated inbox. Thread detection clusters related messages allowing users to collate task information and work on related messages together.

Various approaches have been taken to visualizing inbox threads. These range from combining thread components linearly (Bellotti et al., 2003), to constructing complex tree structures with subthreads (Venolia and Neustaedter, 2003). Most approaches are based on information derived from message subject lines (Rohall et al., 2001, Venolia et al., 2001, Venolia and Neustaedter, 2003, Wattenberg et al., 2005). But as we have seen, topic drift means that subject line is at best a weak indicator of relations between elements of a specific interdependent task. One way to address such topic drift is to allow users to manually intervene, by combining different threads or breaking threads according to their own definition of what is related (Bellotti et al., 2003, 2005). But while these hybrid approaches are promising, most have not as yet been subjected to detailed evaluation, and it may be that to be effective they have to be tuned to an individual user's email or to the user's workgroup.

#### *Embedded Support for Task Management*

Bellotti et al.'s (2003, 2005) prototype TaskMaster provides embedded support for tasks based around the notion of *thrasks*. TaskMaster is designed to present a task-centric view of email addressing the problems of *collation* and *reminding*, introducing the concept of thrasks. The thrask model supports collation as follows: any related incoming messages (replies in a thread, along with associated files or links) are automatically grouped together, using message content and metadata. Each thrask is represented in the inbox by its first message's subject-line, providing a compact representation that serves a reminding function. Clicking on the subject-line reveals the relevant set of messages allowing the user to work on the task. In a field trial involving 9 users, the system performed well with users finding many of the embedded task management features compelling.

Whittaker et al. (1997) also provided direct support for task management in the TeleNotes system. TeleNotes supported interdependent task management by mimicking people's current use of physical desktop piles and spatial location in the following ways. Messages relating to a common task are organized into piles that appear on the users' computer desktop. This serves both to remind users about outstanding tasks, but also helps with collation as related materials are clustered together. The system was positively evaluated by 8 users in an extended field trial.

A slightly different approach to task management was taken in the ContactMap system (Nardi et al., 2002, Whittaker et al., 2004). Instead of organizing the UI in terms of messages involved in a task it organizes in terms of the *people* involved in that task. The UI represents a social network

of the important people in the user's work and social life, and messages (and tasks) can be accessed by clicking on individuals or groups involved in that task. It also supports reminders and alerting, and was positively evaluated in laboratory and field trials.

### *Workflow Systems*

Workflow systems assume that organizational tasks have a predictable structure and that this is associated with different work roles (Prinz, 1993, Winograd, 1994). For example a purchase order may have to be initiated by an employee, approved by a manager, processed by the purchasing department, before being ordered and delivered to the employee. In principle, workflow systems address *interdependent task management*; they support reminding as well as message collation. In practice, however, they have three major limitations: *additional work*, *coverage* and *lack of integration*.

Workflow systems are hard to set up, as workflows are often hard to define, and they often introduce *additional work* for senders. One reason users rejecting early systems such as Information Lens (Malone et al., 1986) and Co-ordinator (Winograd and Flores, et al., 1988) was that these required senders to add additional information to the initiating message (Mackay, 1990). The *coverage* problem arises because many interdependent tasks lack the predictability needed for the workflow approach to succeed. Workflow is effective for tasks that have predictable structure, but as we have seen, most interdependent tasks have an evolving structure and require iterative negotiation for their solution. This makes them inappropriate for workflow tools (Bowers et al., 1995, Dourish et al., 1996, Suchman, 1994). And even when tasks are amenable to workflow, there are *integration* issues. With a few notable exceptions (Borenstein, 1992, Borenstein and Thyberg, 1988), most workflow systems are not well integrated with email clients, so that users have to switch to a separate application - introducing extra cognitive overhead.

## **10.4.4. Organizing Messages into Folders**

*Assisted filing*: Organization is a major problem for email users, and several agent-based systems attempt to provide *assisted filing* (Boone, 1998, Cohen, 1996, Mock, 2001, Segal and Kephart, 1999, Takkinen and Shahmehri, 1998). These systems use machine learning techniques to automatically elicit the defining characteristics of existing folders, based on message headers and content. These definitions can then be used to classify each inbox message suggesting the folder it best matches. Some of these algorithms have been tested offline on message corpora, showing that they can categorize inbox documents with a reasonable degree of success - reporting over 85% accuracy (Cohen, 1996). Although these results are promising, none of the systems has been tested in real usage contexts, and it may be that (as with the filtering techniques described above) users are unwilling to trust systems to assist in filing their messages.

It has been noted that users feel uncomfortable about automated mail classification, since they may not agree with the system's decision (Pazzani, 2000). And if a message is automatically filed, they may be unable to relocate it. MailCat (Segal and Kephart, 1999) improves classification performance and avoids the risk of the user not knowing where content has gone by presenting its best three guesses about where to place the message. Users then actively select



their best choice. Crawford's i-ems intelligent classification application also allows user intervention if the system suggested folder is non-ideal (Crawford et al., 2002).

While these designs address issues of user trust, another serious problem is that *assisted filing* may not cover important aspects of filing. Assisted filing classifies inbox messages into *existing* folder categories, whereas a major user filing problem lies in defining *new* folders (Whittaker and Sidner, 1996). A related problem concerns individual differences in email handling; a subset of users either do not file messages at all, or make rudimentary use of folders (Bälter and Sidner, 2002, Boardman and Sasse, 2004, Whittaker and Sidner, 1996). Assisted filing cannot work for them because they have few folders for the system to learn. A final problem concerns messages that are hard to classify in part because they belong in multiple folders. With some exceptions (e.g. NotesMail) most systems force users to classify messages into a single folder.

## 10.5. Looking Forward: The Future of Email and PIM

We now speculate about the future of email. There are some aspects of email that we do not expect to change in the next few years: (a) *List views*, because they are convenient for viewing, archiving and sorting; (b) *Folders and sorting*, because, even if the system can help search, people will still need multiple ways to organize and find something (c) *Information Overload*, because email continues to be an easy means to distribute one message to many people and ever more collaborative work processes are being moved online (d) *Prevalence of Attachments*, because messages often concern discussion of, and work around other content. We do anticipate changes in email and PIM for our four activities however.

### 10.5.1. Allocating Attention

One recent development that promises to affect email is the popularity of other communication and publishing applications, such as Instant Messaging (IM), Blogs, Wikis, and the Web. A large proportion of email traffic consists of messages that do not necessarily demand a reply (Dabbish et al., 2005). So, instead of using email to distribute such information, it might be better to publish it in a Blog, Wiki or on a webpage, reducing overall email traffic and allowing users to focus their attention on personal emails that demand a response. Similarly, one might migrate urgent messages out of email into more interactive applications such as IM. TeleNotes (Whittaker et al., 1997) does this by integrating email with IM. It presents IM messages on the desktop with alerting, but the same messages can also be viewed from within email if the user chooses. IBM's Activity Explorer (Geyer et al., 2006) is another unified communication application, blurring the boundaries between email chats and shared workspaces, offering integration of six types of objects: message, chat, file, folder, annotated screen shot, and to-do item. Although these developments are promising, there are significant research questions raised by such melded communications. For example, if information is published rather than emailed, how are people alerted about new content? Clumsy implementations of alerts (e.g. sending these through email) replicates the original problem by increasing email volumes. And if we have information distributed in different communication applications this may contribute to fragmentation – making it harder for users to collate related information. On a different note, we

expect that advances in machine learning and text processing will lead to the development of effective methods to prioritize emails, identifying those that are urgent, and allowing users to better focus on important messages.

### **10.5.2. Deciding Actions**

Deciding appropriate message actions is a difficult challenge for automated support, requiring insight into the meaning and importance of email. Given the current state of the art, we hardly want our intelligent email systems deciding to sign us up for distribution lists or to follow sponsored links to purchase products. But recent attention has turned to content-based techniques that identify task-related aspects of messages, so that they can be more easily acted upon by people.

Some of this new work simply focuses on summarizing messages or threads so that it is easier to get the general gist of the content (Abu-Hakima et al., 2001, Muresan et al., 2001, Rambow et al., 2004, Tzoukermann et al., 2001). Other work goes further and attempts to identify the relationship between messages and activities that the user is engaged in. One approach is to recognize messages that are part of a standard automated workflow (such as purchasing a book online) and what steps they represent in such processes (Kushmerik and Lau, 2005). Another is to combine analysis of senders' and recipients' addresses and message bodies to identify messages that are related to a common task, making it easier for the user to monitor it (Dredze et al., 2006). Yet another is to identify the actions expressed in the text of the email body and extract these for the user's optional addition to a to-do list (Corston-Oliver et al., 2004). More advanced systems suggest possible actions on a message. Horvitz et al., (1999) and Metral (1993) built systems that track the history user actions to learn which messages the user reads, deletes, responds to, or files – suggesting appropriate actions for new messages based on these inferred rules. And Carvalho and Cohen (2005) use machine learning to identify messages requiring an action, e.g. requests or commitments.

With continued advances in machine-learning techniques, we expect to see more such systems that assist users in determining how to act upon their email. So far, however, this work is somewhat preliminary and, even where there is a functional email client, problems of trust mean it still remains at users' discretion as to how to act upon the message content. Even as these techniques improve, they will still require both careful integration with existing systems and transparency of operation to engender users' trust in their effectiveness. And one important limitation of all this new work is that little of it has been evaluated with real users.

### **10.5.3. Task Management**

Task management is currently *the* critical unresolved email problem for users, although there are a number of promising trends here. We expect future research will continue to generate novel visualizations of email, to help provide greater structure to the inbox (Kerr, 2003, Viegas et al., 2004, Whittaker et al., 2004). And various novel email clients already provide embedded support for task management, including collation and reminding (Bellotti et al., 2003, 2005, Wattenberg et al., 2005, Whittaker et al., 1997, Whittaker et al., 2004). While only one of these ReMail

(Wattenberg et al., 2005) has been fully integrated into a functioning product, in all cases user trials suggest that they have features that support critical email problems.

What may significantly improve the performance and acceptability of these systems is developments in text processing and machine learning, allowing systems to automatically infer users' tasks (Canny, 2004; Dragunov et al., 2005, Dredge et al., 2006). Other similar techniques will improve inbox organization, by collating related information (Khoussainov and Kushmerick, 2005; Surendran, et al., 2005; Wan and McKeown, 2004). These clustering techniques not only help task management by collating messages relating to outstanding tasks, but they may also clear the inbox by suggesting new folders. While clustering has been applied elsewhere to document collections and web pages (Voorhees, 1996, Willett, 1988, Zamir et al., 1996), as yet it has not been fully evaluated on email. And other research directed towards better message classification will help task management, e.g. by recognizing of messages containing obligations, or standard messages that can be dealt with via autoreply (Carvalho and Cohen, 2005). Of course these developments give rise to important issues that need to be addressed by careful interaction design. They introduce automated processes into a critical application where the cost of algorithmic error without human oversight is high.

#### **10.5.4. Organizing Messages into Folders**

Another recent innovation that promises radical change is the emergence of *desktop search*. Dumais et al., (2003) developed *Stuff I've Seen*, a prototype system that addresses the PIM finding problem by providing search that operates across different applications. This has been followed by commercial desktop search from both Google and Microsoft. Unlike prior desktop or email search, these systems are fast and operate across entire archives retrieving emails, documents, presentations, spreadsheets or documents that relate to the user's query. Desktop search tackles a number of important email problems. It supports *collation* for task management – allowing information items concerning the same task to be retrieved and worked on together. This addresses the significant problem of reuniting originating emails with their attached documents when the documents have been filed in a forgotten location. It should also reduce users' filing burdens by retrieving information regardless of where it has been filed.

While desktop search undoubtedly addresses key email problems it nevertheless suffers from some important limitations. One significant question concerns its accuracy, and information retrieval researchers have suggested that new techniques will be needed to customize existing search tools, so that they can work effectively on private as opposed to public collections (Dumais et al., 2003). Some have speculated that desktop search means that users will no longer file emails, finessing problems of defining and managing folders, but this is unlikely to be the case. Firstly, browsing rather than search remains many users' preferred way to access personal information, even when search is available – in part because of their familiarity with the structure of their own file systems (Barreau and Nardi, 1995, Boardman and Sasse, 2004, Teevan et al., 2004). Folders may also be retained because they provide important semantic information that would be hard to replace by search. For example, folder structure may represent a task decomposition of a project (Jones et al., 2005, Teevan et al., 2004). Most importantly, while search addresses the *finding* problem it does not address other critical aspects of task management, such as reminding oneself that certain tasks remain outstanding. It also cannot help

users to determine which messages they need to allocate attention to, or decide what actions need to be applied to a message.

### 10.5.5. Other Long-Term Research Questions

There are other long-term developments which promise to affect our 4 main email and PIM activities. For example, alternative attentional models could fundamentally change email usage, by imposing payments on sending messages, similar to schemes used by the post office (Kraut et al., 2004). Implementing this would surely reduce email volumes – addressing attention allocation and more indirectly, helping task management. Other key developments may be at the system level, where people have proposed radical solutions to PIM and email problems that demand changes in application architectures, by integrating all PIM activities into email, or by migrating information out of email into dedicated applications (Whittaker et al., 2006). There are important trade-offs here. Centralized systems reduce information fragmentation, but may not provide direct support for key PIM functions. Less radical, but still important, issues concern how we can provide more direct support for the individual differences that are pervasive in PIM and email (see Chapter XX). Finally, despite the centrality of email in our working lives, we lack good theoretical models of asynchronous communication and CMC that might drive new designs, providing new alternatives to our mission-critical but still problematic systems.

## 10.6. References

- Abrams, D., Baecker, R., & Chignell, M. (1998, April 18-23, 1998). *Information archiving with bookmarks: personal web space construction and organization*. Paper presented at the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI 1998), Los Angeles, California.
- Abu-Hakima, S., McFarland, C., and Meech, J. F. 2001. An agent-based system for email highlighting. In Proceedings of the Fifth international Conference on Autonomous Agents (Montreal, Quebec, Canada). AGENTS '01. ACM Press, New York, NY, 224-225.
- Ackerman, M. S. (2000). *Developing for privacy: Civility frameworks and technical design*. Paper presented at the Proc.Computers, Freedom, & Privacy.
- Ackerman, M. S., Cranor, L. F., and Reagle, J. (1999). *Privacy in e-commerce: Examining user scenarios and privacy preferences*. Paper presented at the Proc. ACM Conference on Electronic Commerce.
- Adar, E., Karger, D., & Stein, L. A. (1999, November 2-6). *Haystack: per-user information environment*. Paper presented at the 8th Conference on Information and Knowledge Management (CIKM 1999), Kansas City, MO.
- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Bälter, O. (2000, April 1-6). *Keystroke level analysis of email message organization*. Paper presented at the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI 2000), Hague, Netherlands.

- Bälter, O., & Sidner, C. L. (2002). *Bifrost inbox organizer: giving users control over the inbox*. Paper presented at the Second Nordic Conference on Human-computer interaction.
- Barreau, D. K. (1995). Context as a factor in personal information management systems. *Journal of the American Society for Information Science*, 46(5), 327-339.
- Barreau, D. K., & Nardi, B. A. (1995). Finding and reminding: file organization from the desktop. *SIGCHI Bulletin*, 27(3), 39-43.
- Barsalou, L. (1991). Deriving categories to achieve goals. *The Psychology of Learning and Motivation*, 27, 1-64.
- Bellotti, V., Dalal, B, Good, N, Flynn, P, Bobrow, D. G., & Ducheneaut, N. (2004). *What a To-Do: Studies of Task Management Towards the Design of a Personal Task List Manager*. Paper presented at the Conference on Human Factors in Computing Systems, Vienna, Austria.
- Bellotti, V., Ducheneaut, N., Howard, M., Neuwirth, C., & Smith, I. (2002, June 25-28). *Innovation in extremis : evolving an application for the critical work of email and information management*. Paper presented at the Conference on Designing Interactive Systems (DIS2002), London.
- Bellotti, V., Ducheneaut, N., Howard, M., & Smith, I. (2003). *Taking email to task: The design and evaluation of a task management centered email tool*. Paper presented at the Conference on Human Factors in Computing Systems (CHI 2003), Ft. Lauderdale, Florida.
- Bellotti, V., & Smith, I. (2000, Aug 2000). *Informing the design of an information management system with iterative fieldwork*. Paper presented at the Conference on Designing interactive systems (DIS 2000), New York City, New York.
- Canny, J. (2004). GaP: A Factor Model for Discrete Data. in Proceedings of SIGIR. ACM Press, 2004.
- Corston-Oliver, S., E. Ringger, M. Gamon and R. Campbell. 2004. Task-focused summarization of email. Proceedings of the Text Summarization Branches Out workshop, ACL 2004.
- Crawford, E., Kay, J., and McCreath, E. 2002. An intelligent interface for sorting electronic mail. In Proceedings of the 7th international Conference on intelligent User interfaces (San Francisco, California, USA, January 13 - 16, 2002). IUI '02. ACM Press, New York, NY, 182-183
- Dredze, M., Lau, T, and Kushmerick, N (2006). Automatically Classifying Emails into Activities, IUI 2006, Sydney, Australia, January 2006.
- Dragunov, A., Dietterich, T., Johnsrude, K., McLaughlin, M., Li, L., Herlocker, J. (2005). TaskTracer: A Desktop Environment to Support Multi-tasking Knowledge Workers. International Conference on Intelligent User Interfaces (IUI-2005), (pp. 75-82). ACM Press.
- Ducheneaut, N. and Bellotti, V. (2003). "Ceci n'est pas un objet? Talking about things in email." *Human-Computer Interaction*, 18 (1 & 2), pp. 85-110.
- Flores, F., Graves, M., Hartfield, B., & Winograd, T. (1988). Computer systems and the design of organizational interaction. *ACM Transactions on Office Information Systems*, 6, 153-172.
- Gwizdka, J. (2004a). *Email Task Management Styles: The Cleaners and the Keepers*. Extended Abstracts of ACM Conference on Human Factors in Computing Systems CHI'2004. ACM Press. 1235 - 1238.
- Gwizdka, J. (2004b). *Cognitive Abilities and Email Interaction: Impacts of Interface and Task*. Unpublished PhD Dissertation. University of Toronto.
- Gwizdka, J., Chignell, M.H. (2004). Individual Differences and Task-based User Interface Evaluation: A Case Study of Pending Tasks in Email. *Interacting with Computers*. Elsevier Science. 16 (4). 769-797.

- Horvitz, E., Jacobs, A., Hovel, D. (1999). Attention-Sensitive Alerting, Proceedings of UAI '99, Conference on Uncertainty and Artificial Intelligence, July 1999, Morgan Kaufmann Publishers: San Francisco. pp. 305-313.
- Kerr, Bernard (2003). Thread Arcs: An Email Thread Visualization. 2003 IEEE Symposium on Information Visualization, pp. 27.
- Khoussainov, R. and Kushmerick, N. (2005) Email Task Management: An Iterative Relational Learning Approach. Paper presented at the 2nd International Conference on Email and Anti-Spam (CEAS 2005), Stanford University, CA, USA, July 21-22, 2005.
- Kiritchenko, S. and Matwin, S (2001). Email classification with co-training. In CASCON '01: Proceedings of the 2001 conference of the Centre for Advanced Studies on Collaborative research, pages 192–201. IBM Press, 2001.
- Klimt, B., & Yang, Y. (2004). Introducing the Enron Corpus. First Conference on Email and Anti-Spam (CEAS 2004), Mountain View, CA. July 30<sup>th</sup>-31<sup>st</sup>, 2004.
- Kraut, Robert E., Sunder, Shyam, Morris, James H., Telang, Rahul, Filer, Darrin and Cronin, Matthew A., "Markets for Attention: Will Postage for Email Help?" Human Computer Interaction, 2005.
- Kushmerick, N and Lau, T. (2005) Automated Email Activity Management: An Unsupervised Learning Approach, IUI 2005, San Diego, CA, January 2005.
- Mackay, W. E. (1988). Diversity in the use of electronic mail: A preliminary inquiry. ACM Transactions on Information Systems, 6, 380-397.
- Malone, T. W., Grant, K. R., and Turbak, F. A. (1986). The information lens: an intelligent system for information sharing in organizations. Paper presented at the Conference on Human Factors in Computing Systems (CHI '86). ACM Press, New York, NY, 1-8.
- Metral, M. (1993). Design of a generic learning interface agent. Bachelor of Science thesis, Massachusetts Institute of Technology, Cambridge MA.
- Muresan, S. Tzoukermann, E. Klavans, J.L. (2001). Combining Linguistic and Machine Learning Techniques for Email Summarization. Proceedings of the CoNLL 2001 Workshop at the ACL/EACL 2001
- Pazzani, M. J. (2000). Representation of electronic mail filtering profiles: A user study. Proceedings of the IUI 2000 International Conference on Intelligent User Interfaces. New York: ACM.
- Rambow, O. Shrestha, L. Chen, J. and Laurdisen, C. (2004) Summarizing Email Threads. Proceedings of HLT-NAACL 2004: Short Papers, 2004 - acl.ldc.upenn.edu
- Surendran, A. C., Platt, J. C., & Renshaw, E. (2005). Automatic Discovery of Personal Topics to Organize Email. Paper presented at the 2nd Conference on Email and Anti-Spam (CEAS'05), Stanford University, USA.
- Tzoukermann, E., Muresan, S. and Klavans, J.L. (2001) GIST-IT: Summarizing Email using Linguistic Knowledge and Machine Learning. Proceedings of HLT/KM 2001 Workshop at ACL/EACL 2001 Conference, Toulouse, France, July 2001.
- Viegas, F., boyd, d., Nguyen, D., Potter, J., Donath, J. (2004). Digital Artifacts for Remembering and Storytelling: PostHistory and Social Network Fragments. Proceedings of the 37th Hawaii International Conference on System Sciences, January 2004, pp. 40109a.
- Wan, S. and McKeown, K. (2004) Generating Overview Summaries of Ongoing Email Thread Discussions. In Proceedings of COLING 2004, the 20th International Conference on Computational Linguistics. Geneva, Switzerland.