

NSF Workshop on PIM: Problems and Opportunities “From PIM to GIM” Breakout Group Report

**January 2005
University of Washington**

The “From PIM to GIM” Breakout Group

Facilitators: Thomas Erickson, Jonathan Grudin

Participants: Sue Dumais, Alon Halevy, Kent Unruh, Steve Whittaker

1. Preface

The charter of this group was to explore the question of how personal information management (PIM) relates to group information management (GIM). Motivating questions included how and where PIM fits in the business world, what happens to PIM when the information is not personally owned, and the implications of sharing personal information (intentionally or unintentionally).

As GIM does not exist as a recognized subfield or even phrase, the discussion began with an effort to agree upon a definition. The definitional effort quickly turned into an interplay between attempts to define a model of GIM, and the discussion of particular examples of group information management. As the group converged on a definition of PIM, attention shifted to a consideration of the problems and opportunities offered by GIM. Here, as well, there was an interplay between the discussion of examples and the articulation of problems and opportunities.

2. From PIM to GIM

Broadly put, personal information management (hereafter PIM) serves two ends, instrumental and symbolic. First, artifacts such as “to do” lists, calendars and rolodexes serve as external memories, and enable their users to efficiently conduct their daily tasks. Second, PIM can assist users in managing the impressions that others form of them. Thus, the use of a Day-Timer® or other personal organizer system—available in a wide array of materials (“an expression of your unique style,” according to the Day-Timer web site)—can contribute to creating the impression of the user as an productive, well organized professional. And of course, if the PIM artifacts are adroitly deployed and succeed in achieving their instrumental ends, the impression will be augmented by efficiency of the user’s performance.

PIM can also be seen as functioning in two spheres: private and more public. In the private sphere PIM simply supports one’s personal tasks. Thus, while one’s PIM activities may be glimpsed by others—as when we see someone checking an item off a list or looking up a number in a rolodex—the information is purely for the use of its owner. But PIM also functions in a more public sphere. That is, information is often created with some degree of sharing in mind. A student may take notes, writing a bit more carefully than usual, to share with an absent friend. Or members of a workgroup may develop a practice of sharing their calendars with one another to facilitate meeting scheduling. While this sharing serves the instrumental ends that motivated it, the information thus shared also becomes grist for possible inferences about the owner: the student’s handwriting may be sloppy or her notes incomplete; a workgroup member’s calendar may reveal private information such as medical appointments, or consistently long lunch dates. Thus, when personal information is shared, it introduces tensions between the instrumental ends for which it is shared and the not necessarily desirable inferences that it may support.

The tensions that occur as personal information is shared are complex and intertwined, and moreover have the potential to feed back and alter norms having to do with what is shared, and how it is shared. As a consequence, this area seems a valuable focus for research attention. We have adopted the phrase “Group Information Management” (hereafter GIM) to refer to PIM as it functions in more public spheres. More specifically, we define it as follows: *GIM has to do with how personal information is shared amongst a group, with an emphasis on the norms that underlie that sharing, and the ways in which participants negotiate those norms in response to a variety of tensions.*

3. Examples of GIM

Because GIM has to do with how information is shared amongst a group, it is not surprising that a wide array of applications can be used to support GIM, including email, web pages, WIKIs, and traditionally produced documents. However, while many applications can be turned to GIM ends, there are some that fall more squarely into the GIM arena.

Shared Calendars

An early example of GIM in the digital realm is the development of online calendaring systems. In the 1980's, various developers produced digital analogs of personal calendars that were designed to be shared by groups in an organization. The motivating idea was quite simple: making a person's calendar available to others could facilitate the sometimes onerous task of scheduling a meeting. However, these electronic calendaring systems encountered resistance for reasons ranging from the fact that personal calendars rarely contained a complete picture of their users' availability, to users' realization that calendar information could be used for other not necessarily desirable ends, such as making inferences about users' personal activities. As shared calendaring systems have been adopted and 'naturalized' within organizations, a variety of technical features and social practices have arisen in response to such tensions.

Blogs

A more recent application genre is the blog, a web-based, person-centric diary-like document consisting of relatively short entries displayed in reverse chronological order; these entries can be linked to, and commented upon, by readers. Most blogs are published by individuals for small audiences comprised of family and friends; however, some are published by groups for the explicit purpose of sharing information and generating commentary from a larger audience. Blogs raise interesting issues about audience: Who does the blog author imagine that he or she is writing for? What are the consequences when personal information published in a blog receives attention from a different audience than intended? What steps do blog authors take to avoid these consequences or recover from them?

Social Networking Services

Social networking services such as Orkut, LinkedIn® and Friendster® allow their users to post personal profiles, pictures, and create links to others signifying professional or social ties. The networks of links thus formed can then be viewed, traversed and used to distribute messages. Such systems serve a variety of purposes from supporting online professional networking to enabling singles to find prospective dates. Social networking services raise interesting issues about what users choose to reveal or conceal, how their disclosure of personal information is related to the ends that they hope to achieve, and the ethics of 'counterfeiting' links or conspiring to garner 'inauthentic' recommendations to increase their stature in the system.

Electronic Medical Records

As the information technology systems of the medical and insurance industries become increasingly interlinked, electronic medical records become an increasingly interesting example of GIM. Any particular patient's medical record is composed of information generated by multiple people (and devices); those who contribute to the record may come from different institutions, and enter information for a variety of different purposes. Access to records is by a similarly disparate audience for even more diverse purposes, and questions of ownership and access privileges are complex. This application of GIM raises complex questions of privacy and access and of ownership.

This list is not, of course, an exhaustive one. Other GIM-centered application areas include peer to peer file sharing, information sharing and tagging systems such as del.ici.ous and flickr, online reviewing and rating systems, and event organizing applications such as eVite® and MeetUp™.

4. Issues and Opportunities in GIM

After working towards a definition of GIM, and generating a list of core examples of GIM, the breakout group generated a list of research issues and opportunities that arise in GIM.

4.1. A Simple Model

To organize these issues, let's start with an overly simple model of GIM, and examine each part of the model:

- A person generates information...
- ... that is shared with a group...
- ... in support of some task

Thus, an employee enters appointments in her calendar to share with her coworkers to facilitate the scheduling of meetings. Or a person creates a profile in an online social networking system to be shared with other members of the system for the purpose of getting dates.

A person generates information...

One set of GIM issues has to do with the creation of the to be shared information. What information do people choose to share, and why? (The implication of the model that they do so in support of a task, with a particular group in mind, is simply a conjecture, and, regardless, doesn't explain all cases.) What are the psychological issues that attend the decision to share information (for example, people have been observed to 'clean up' information before sharing it)? What are the various norms that attend sharing, and how do they vary according to form, content and domain of the information?

...that is shared with a group...

Another set of issues has to do with whom the information is shared. How is the audience for the information specified? How does the imagined audience interact with the nature of the information shared? What are the consequences of changes in the audience over time (for example, as an organizationally defined group changes composition)? What are the consequences of 'leakage' of the information beyond the intended audience? And, to the extent that GIM users are concerned with such questions, how might GIM systems support them in preventing or mitigating these issues?

...in support of some task

If we accept that in some cases people choose to share information in support of an envisioned task, what happens when the information turns out to be useful for other tasks that are not in the user's best interests? To what extent is it possible to give users control over uses of their personal information? To what extent is it possible to allow them to retract it after the intended task is completed? To what extent is it possible to simply allow users to be aware of when their information is actually used?

4.2. More Complex Variants of the Model

The above model is quite simple. Let's consider a few variations on it.

Ownership is complicated

The simple model assumes that personal information is owned by an individual. But in fact this assumption can fail in many ways. It may be that the individual is not voluntarily generating the information, as implied in the simple model, but is generating information as a side effect of his or her activities (e.g. credit information; medical records; calling records). In such cases individuals may not own their information, or if they own it may nevertheless lack complete control over its content, distribution or use. This raises a host of issues about who can see the information, how it can be used, whether it can be corrected if in error, or retracted if no longer needed, and how to deal with real or asserted errors in the information, or its distribution or use.

Information is generated collectively

In some cases groups may generate information collectively, creating either a single collective product (e.g. the contents of a Wiki), or a set of individual products that are shared with the other members of the group (e.g. the profiles in a social networking system). These cases are interesting because group norms and incentive systems come into play, and their establishment, support and evolution can play a critical role in shaping the character of the system. These effect the system at all levels, including the nature of the information shared, and how collectively produced information is structured and the maintained over time.

GIM occurs in the context of an institution

GIM often occurs in an institutional context, and thus is shaped by institutional values, practices and mechanisms. Examining the practices of institutions that have developed expertise at GIM (e.g. the Mayo Clinic, with its century of experience in maintaining, sharing and glossing patient records) seems one fruitful avenue of exploration. Similarly, looking at the ways in which the needs of differing institutions (e.g. the medical and insurance industries) shape the nature and use of personal information also seems of interest.

5. Summary

GIM does not currently exist as a distinct field. Even the term, which occasionally occurs in the literature, mostly in the company of PIM, does not have an agreed upon meaning. As the discussion outlined in this document makes clear, GIM raises a number of interesting issues, and has considerable potential as a focus for research.